HBR'S 10 MUST READS

BONUS ARTICLE
"Begin with Trust"
By Frances X. Frei and Anne Morriss

The definitive management ideas of the year from **Harvard Business Review.**

2022

5 Years of Must Reads from HBR: 2022 Edition

Contents

Harvard Business Review Press
Boston, Massachusetts

eISBN: 978-1-6478-2440-2

HBR'S 10 MUST READS

BONUS ARTICLE

"Customer Loyalty Is Overrated"

By A.G. Lafley and Roger L. Martin

The definitive management ideas of the year from **Harvard Business Review.**

2018

HBR's 10 Must Reads series is the definitive collection of ideas and best practices for aspiring and experienced leaders alike. These books offer essential reading selected from the pages of *Harvard Business Review* on topics critical to the success of every manager.

Titles include:

HBR's 10 Must Reads 2015
HBR's 10 Must Reads 2016
HBR's 10 Must Reads 2017
HBR's 10 Must Reads 2018
HBR's 10 Must Reads for New Managers
HBR's 10 Must Reads on Change Management
HBR's 10 Must Reads on Collaboration
HBR's 10 Must Reads on Communication
HBR's 10 Must Reads on Emotional Intelligence
HBR's 10 Must Reads on Innovation
HBR's 10 Must Reads on Leadership
HBR's 10 Must Reads on Making Smart Decisions
HBR's 10 Must Reads on Managing Across Cultures
HBR's 10 Must Reads on Managing People
HBR's 10 Must Reads on Managing Yourself
HBR's 10 Must Reads on Sales
HBR's 10 Must Reads on Strategic Marketing
HBR's 10 Must Reads on Strategy
HBR's 10 Must Reads on Teams

HBR’s 10 Must Reads: The Essentials

HBR'S 10 MUST READS

The definitive management ideas of the year from

Harvard Business Review.

2018

HARVARD BUSINESS REVIEW PRESS

Boston, Massachusetts

Copyright

HBR Press Quantity Sales Discounts

Harvard Business Review Press titles are available at significant quantity discounts when purchased in bulk for client gifts, sales promotions, and premiums. Special editions, including books with corporate logos, customized covers, and letters from the company or CEO printed in the front matter, as well as excerpts of existing books, can also be created in large quantities for special needs.

For details and discount information for both print and ebook formats, contact booksales@harvardbusiness.org, tel. 800-988-0886, or www.hbr.org/bulksales.

Copyright 2018 Harvard Business School Publishing Corporation
All rights reserved

No part of this publication may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form, or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the prior permission of the publisher. Requests for permission should be directed to permissions@hbsp.harvard.edu, or mailed to Permissions, Harvard Business School Publishing, 60 Harvard Way, Boston, Massachusetts 02163.

First eBook Edition: Oct 2017

ISBN: 9781633693067

eISBN: 9781633693074

Contents

Editors' Note

Every year, as we build each issue of *Harvard Business Review*, we examine the most important challenges facing business leaders today, from technology to people management. Rather than simply monitoring buzzwords or headlines, this involves a combination of looking forward to how businesses will need to incorporate new technologies and contextual realities, and also looking back at lingering management problems to find the ways that researchers and practitioners are addressing them today. The standout articles of the year collected here, for example, explain emerging phenomena like blockchain, dataviz literacy, and algorithms in practical terms. They also offer new perspectives on long-term issues such as boosting employee engagement, increasing diversity, and fixing the U.S. health care system. We showcase these and other critical themes highlighted by our authors from the past year of *Harvard Business Review* in this volume.

In today's crowded and competitive marketplace, companies often feel pressure to rebrand or expand their offerings to stay alive. But P&G's A.G. Lafley and strategy expert and Rotman School of Management professor Roger L. Martin say companies should focus their efforts on strengthening customers' habits, not developing products or redesigning packaging. In **"Customer Loyalty Is Overrated,"** the authors

acknowledge that although it’s hard work to establish a brand, once you’ve done so, constant reinvention won’t keep customers coming back. Research suggests that what makes competitive advantage sustainable is helping consumers avoid expending the mental energy to make a choice. Customers don’t want to have to evaluate their options every time they shop; they just want to buy what they’ve always bought. And each time customers pick the same product, they boost its advantage over that of the products they didn’t choose.

Inconsistent decision making is often a hidden and expensive problem plaguing companies—not the big, sweeping, strategy-related choices, but the daily decisions and judgment calls, which can swing radically from one individual to the next. This problem affects not just new employees but seasoned people who have been in the same roles, following the same well-established guidelines. Irrelevant factors, such as mood and the weather, can affect a person’s decisions from one occasion to the next. This chance variability of decisions is called *noise*. In **“Noise: How to Overcome the High, Hidden Cost of Inconsistent Decision Making,”** Nobel laureate and Princeton psychology professor Daniel Kahneman and data analysis experts Andrew M. Rosenfield, Linnea Gandhi, and Tom Blaser explain how organizations can perform a “noise audit” and use algorithms and simple commonsense rules to guide employees toward making more-consistent decisions.

Managers should all be relying more on data in their decision making, but it arrives at such velocity, and in such volume, that many of them don't know quite what to do with it. A good first step is to create a visualization or a chart. To do that well, however, you need to understand the nature of your data and keep your purpose in mind, according to Scott Berinato, an HBR senior editor and the author of *Good Charts: The HBR Guide to Making Smarter, More Persuasive Data Visualizations*. That strategic attitude will make your charts and presentations much clearer and more effective. In **"Visualizations That Really Work,"** Berinato outlines categories of approach and the tools and resources you'll need for each.

Managers are pretty good at assessing *whether* a new technology will overtake an existing one, but they haven't quite figured out how to know *when* that will happen. In **"Right Tech, Wrong Time,"** professors Ron Adner and Rahul Kapoor say that not just your new technology but also the ecosystem in which it will exist—the related technologies, services, standards, and regulations—can influence how quickly it's adopted. They provide a framework to assess how soon disruptive change is coming to your industry by analyzing the dynamics of the context in which it will exist. If the new technology doesn't need a new ecosystem to support it—if it's essentially plug-and-play—adoption will be swift. But if complements are needed (for example, electric cars require a

network of charging stations), the pace of substitution will slow until those challenges have been resolved.

How to pay for health care is a problem the United States has struggled with for a long time. Fee-for-service, the dominant model today, is widely recognized as the single biggest obstacle to improving health care delivery, because it rewards the quantity rather than the quality or efficiency of care. What we need is a system that rewards providers for delivering superior value to patients—for achieving better health outcomes at a lower cost. In **"How to Pay for Health Care,"** strategy giants Michael E. Porter and Robert S. Kaplan argue that a "bundled payments" model is the right one, because it triggers competition among providers to create value where it matters—at the individual patient level. They describe robust proof-of-concept initiatives in the United States and abroad that show how the challenges of transitioning to bundled payments are already being overcome.

Another system that's overdue for reform is annual performance reviews. Emphasizing individual accountability for past results, traditional appraisals give short shrift to improving current performance and developing talent for the future. That can hinder long-term competitiveness, say Peter Cappelli and Anna Tavis in **"The Performance Management Revolution."** To better support employee development, many organizations are dropping or radically changing their annual-review systems in favor of giving people less-formal, more-

frequent feedback that follows the natural cycle of work. The authors explain how performance management has evolved over the decades and why current thinking has shifted.

Goal-setting and evaluation are one way to motivate your employees, but how to engage them is another long-standing issue for managers and organizations. Francesca Gino, a professor of business administration at Harvard Business School, conducted groundbreaking research and found that whether consciously or unconsciously, organizations pressure employees—including leaders—to reserve their real, authentic, nonconforming selves for outside the workplace. This pressure to conform, she writes in **"Let Your Workers Rebel,"** can have a significant negative impact on engagement, productivity, and the ability to innovate. To fix this problem, she says, develop a culture that supports "constructive nonconformity": encourage your workers to break rules and be themselves.

Diversity programs are another relic in organizations: Most companies rely on the same approach they've been using since the 1960s to reduce bias and increase diversity—one that focuses on controlling managers' behaviors. But as studies have shown, that tends to activate bias rather than quash it, because people rebel against rules that threaten their autonomy. In the McKinsey Award–winning **"Why Diversity Programs Fail,"** Frank Dobbin and Alexandra Kalev draw on their research to suggest ways of promoting diversity that engage employees in working explicitly toward that goal, increase contact with female and

minority colleagues to lessen bias, and encourage social accountability through transparency and diversity task forces.

The U.S. presidential election in November 2016 left in its wake a question that also resonates in other countries experiencing populist upwellings: How did the liberal political establishment, media, and electorate fail to anticipate the anger and desperate desire for change that ushered in the Trump administration? In **"What So Many People Don't Get About the U.S. Working Class,"** Joan C. Williams, a distinguished professor of law at UC Hastings, points her finger at "class cluelessness" and draws on her expertise in labor and social class to describe to "professional elites" the difference between "working-class" and poor, the role of the urban-rural divide, the need for job and college programs, and how race and gender do (and don't!) play a part in working-class politics.

We've all heard that blockchain will revolutionize business. But what is it? And when will organizations need to integrate it into their daily operations? In **"The Truth About Blockchain,"** Marco Iansiti and Karim R. Lakhani, academics who study digital innovation in business, explain this new technology and assure us that its arrival is going to take a lot longer than many people claim. Like TCP/IP (on which the internet was built), blockchain is a foundational technology that will require broad coordination. Its level of complexity—technological, regulatory, and social—will be unprecedented. It could transform the economy by slashing the cost of transactions (and how long

they take) and eliminating intermediaries such as lawyers and bankers. The adoption of TCP/IP suggests that blockchain will follow a fairly predictable path. But although the journey may take years, it's not too early to start planning.

New technology is born of effective R&D, but numerous potential stumbling blocks lie between research and commercial development. Early-stage research is expensive, risky, and unpredictable—so corporations generally shy away from it, leaving many opportunities unexplored. They could revitalize their research operations by adopting the approach taken by Bob Langer, a chemical engineer whose lab at MIT is one of the most productive and profitable research facilities in the world. **"The Edison of Medicine,"** by HBR senior editor Steven Prokesch, details Langer Lab's proven formula for accelerating the pace of discoveries and getting them into the world as products. It includes focusing on projects that could make the most difference to society, finding opportunity in the constant turnover of researchers, and cultivating a leadership style that balances freedom and support.

Looking across disciplines and trends and synthesizing the best ideas is important—and time-consuming—work for today's leaders. With this volume, we've done some of that heavy lifting for you. With topics ranging from a new type of literacy to a new way to record transactions, the articles here will help you better manage your work today and make smart plans for whatever lies ahead.

—The Editors

Customer Loyalty Is Overrated

by A.G. Lafley and Roger L. Martin

LATE IN THE SPRING OF 2016 Facebook's category-leading photo-sharing application, Instagram, abandoned its original icon, a retro camera familiar to the app's 400-million-plus users, and replaced it with a flat modernist design that, as the head of design explained, "suggests a camera." At a time when Instagram was under a growing threat from its rival Snapchat, he offered this rationale for the switch: The icon "was beginning to feel … not reflective of the community, and we thought we could make it better."

The assessment of *AdWeek,* the marketing industry bible, was clear from its headline: "Instagram's New Logo Is a Travesty. Can We Change it Back? Please?" In *GQ*'s article "Logo Change No One Wanted Just Came to Instagram," the magazine's panel of designers called the new icon "honestly horrible," "so ugly," and "trash," and summarized the change thus: "Instagram spent YEARS building up visual brand equity with its existing logo, training users where to tap, and now instead of iterating

on that, it's flushing it all down the toilet for the homescreen equivalent of a Starburst."

It's too soon to tell whether the design change will actually have commercial consequences for Instagram, but this is not the first time a company has experienced such a reaction to a rebranding or a relaunch. PepsiCo's introduction of its aspartame-free Diet Pepsi was—like the infamous New Coke debacle—a botched attempt at reinvention that resulted in serious revenue losses and had to be reversed. The interesting question, therefore, is: Why do well-performing companies routinely succumb to the lure of radical rebranding? One could understand the temptation to adopt such a strategy in the face of disaster, but Instagram, PepsiCo, and Coke were hardly staring into the abyss. (It's worth noting that Snapchat, whose market share among young users is now particularly strong, has assiduously stuck to its familiar ghost icon. Full disclosure: A.G. Lafley serves on the board of Snap Inc.)

The answer, we believe, is rooted in some serious misperceptions about the nature of competitive advantage. Much new thinking in strategy argues that the fast pace of change in modern business (perhaps nowhere more obvious than in the app world) means no competitive advantage is sustainable, so companies must continually update their business models, strategies, and communications to respond in real time to the explosion of choice that ever more sophisticated consumers now face. To keep your customers—

and to attract new ones—you need to remain relevant and superior. Hence Instagram was doing exactly what it was supposed to do: changing proactively.

That's an edgy thought, to be sure; but a lot of evidence contradicts it. Consider Southwest Airlines, Vanguard, and IKEA, all featured in Michael Porter's classic 1996 HBR article "What Is Strategy?" as exemplars of long-lived competitive advantage. A full two decades later those companies are still at the top of their respective industries, pursuing largely unchanged strategies and branding. And although Google, Facebook, or Amazon might stumble and be crushed by some upstart, the competitive positions of those giants hardly look fleeting. Closer to home (one author of this article is part of the P&G family), it would strike the Tide or Head & Shoulders brand managers of the past 50 years as rather odd to hear that their half-century advantages have not been or are not sustainable. (No doubt the Unilever managers of long-standing consumer favorites such as Dove soap and Hellmann's mayonnaise would feel the same.)

In this article we draw on modern behavioral research to offer a theory about what makes competitive advantage last. It explains both missteps like Instagram's and success stories like Tide's. We argue that performance is sustained not by offering customers the perfect choice but by offering them the easy one. So even if a value proposition is what first attracted them, it is not necessarily what keeps them coming.

Idea in Brief

The Problem

Product innovations often flame out on launch, despite tremendous efforts to make them attractive, relevant, and up-to-date.

Why It Happens

Customers don't want to spend the mental energy needed to choose between products.

The Solution

To strengthen customers' habits, innovations should represent a progression of the brand rather than a break with the past.

In this alternative worldview, holding on to customers is not a matter of continually adapting to changing needs in order to remain the rational or emotional best fit. It's about helping customers avoid having to make yet another choice. To do that, you have to create what we call *cumulative advantage*.

Let's begin by exploring what our brains actually do when we shop.

Creatures of Habit

The conventional wisdom about competitive advantage is that successful companies pick a position, target a set of consumers,

and configure activities to serve them better. The goal is to make customers repeat their purchases by matching the value proposition to their needs. By fending off competitors through ever-evolving uniqueness and personalization, the company can achieve sustainable competitive advantage.

An assumption implicit in that definition is that consumers are making deliberate, perhaps even rational, decisions. Their reasons for buying products and services may be emotional, but they always result from somewhat conscious logic. Therefore a good strategy figures out and responds to that logic.

But the idea that purchase decisions arise from conscious choice flies in the face of much research in behavioral psychology. The brain, it turns out, is not so much an analytical machine as a gap-filling machine: It takes noisy, incomplete information from the world and quickly fills in the missing pieces on the basis of past experience. Intuition—thoughts, opinions, and preferences that come to mind quickly and without reflection but are strong enough to act on—is the product of this process. It's not just what gets filled in that determines our intuitive judgments, however. They are heavily influenced by the speed and ease of the filling-in process itself, a phenomenon psychologists call *processing fluency.* When we describe making a decision because it "just feels right," the processing leading to the decision has been fluent.

Processing fluency is itself the product of repeated experience, and it increases relentlessly with the number of times we have the experience. Prior exposure to an object improves the ability to perceive and identify that object. As an object is presented repeatedly, the neurons that code features not essential for recognizing the object dampen their responses, and the neural network becomes more selective and efficient at object identification. In other words, repeated stimuli have lower perceptual-identification thresholds, require less attention to be noticed, and are faster and more accurately named or read. What's more, consumers tend to prefer them to new stimuli.

In short, research into the workings of the human brain suggests that the mind loves automaticity more than just about anything else—certainly more than engaging in conscious consideration. Given a choice, it would like to do the same things over and over again. If the mind develops a view over time that Tide gets clothes cleaner, and Tide is available and accessible on the store shelf or the web page, the easy, familiar thing to do is to buy Tide yet another time.

A driving reason to choose the leading product in the market, therefore, is simply that it is the easiest thing to do: In whatever distribution channel you shop, it will be the most prominent offering. In the supermarket, the mass merchandiser, or the drugstore, it will dominate the shelf. In addition, you have probably bought it before from that very

shelf. Doing so again is the easiest possible action you can take. Not only that, but every time you buy another unit of the brand in question, you make it easier to do—for which the mind applauds you.

Meanwhile, it becomes ever so slightly harder to buy the products you didn't choose, and that gap widens with every purchase—as long, of course, as the chosen product consistently fulfills your expectations. This logic holds as much in the new economy as in the old. If you make Facebook your home page, every aspect of that page will be totally familiar to you, and the impact will be as powerful as facing a wall of Tide in a store—or more so.

Buying the biggest, easiest brand creates a cycle in which share leadership is continually increased over time. Each time you select and use a given product or service, its advantage over the products or services you didn't choose cumulates.

The growth of cumulative advantage—absent changes that force conscious reappraisal—is nearly inexorable. Thirty years ago Tide enjoyed a small lead of 33% to 28% over Unilever's Surf in the lucrative U.S. laundry detergent market. Consumers at the time slowly but surely formed habits that put Tide further ahead of Surf. Every year, the habit differential increased and the share gap widened. In 2008 Unilever exited the business and sold its brands to what was then a private-label detergent manufacturer. Now Tide enjoys a greater than 40% market share, making it the runaway leader in the U.S.

detergent market. Its largest branded competitor has a share of less than 10%. (For a discussion of why small brands even survive in this environment, see the sidebar "The Perverse Upside of Customer Disloyalty.")

A Complement to Choice

We don't claim that consumer choice is never conscious, or that the quality of a value proposition is irrelevant. To the contrary: People must have a reason to buy a product in the first place. And sometimes a new technology or a new regulation enables a company to radically lower a product's price or to offer new features or a wholly new solution to a customer need in a way that demands consumers' consideration.

The Perverse Upside of Customer Disloyalty

IF CONSUMERS ARE SLAVES OF HABIT, it's hard to argue that they are "loyal" customers in the sense that they consciously attach themselves to a brand on the assumption that it meets rational or emotional needs. In fact, customers are much more fickle than many marketers assume: Often the brands that are believed to depend on loyal customers achieve the lowest loyalty scores.

For example, Colgate and Crest are the leading toothpaste brands in the U.S. market, with about 75% of it between them.

Customers for both are loyal 50% of the time (their preferred brand accounts for 50% of their annual toothpaste purchases). Tom's toothpaste, a niche "natural" brand based in Maine, has a 1% market share and is thought to have a fanatical customer following. One might expect the data to show that the 1% are mostly repeat buyers. But in fact Tom's customers are loyal only 25% of the time—half the rate of the big brands.

So why do fringe brands like Tom's survive? The answer, perhaps perversely, is that with big-brand loyalty rates at 50%, just enough customers will buy small brands from time to time to keep the latter in business. But the small brands can't overcome the familiarity barrier, and although entirely new brands do enter categories and become leaders, it is extremely rare for a small fringe brand to successfully take on an established leader.

Robust where-to-play and how-to-win choices, therefore, are still essential to strategy. Without a value proposition superior to those of other companies that are attempting to appeal to the same customers, a company has nothing to build on.

But if it is to extend that initial competitive advantage, the company must invest in turning its proposition into a habit rather than a choice. Hence we can formally define cumulative advantage as the layer that a company builds on its initial competitive advantage by making its product or service an ever more instinctively comfortable choice for the customer.

Companies that don’t build cumulative advantage are likely to be overtaken by competitors that succeed in doing so. A good example is Myspace, whose failure is often cited as proof that competitive advantage is inherently unsustainable. Our interpretation is somewhat different.

Launched in August 2003, Myspace became America’s number one social networking site within two years and in 2006 overtook Google to become the most visited site of any kind in the United States. Nevertheless, a mere two years later it was outstripped by Facebook, which demolished it competitively—to the extent that Myspace was sold in 2011 for $35 million, a fraction of the $580 million that News Corp had paid for it in 2005.

Why did Myspace fail? Our answer is that it didn’t even try to achieve cumulative advantage. To begin with, it allowed users to create web pages that expressed their own personal style, so individual pages looked very different to visitors. It also placed advertising in jarring ways—and included ads for indecent services, which riled regulators. When News Corp bought Myspace, it ramped up ad density, further cluttering the site. To entice more users, Myspace rolled out what *Bloomberg Businessweek* referred to as “a dizzying number of features: communication tools such as instant messaging, a classifieds program, a video player, a music player, a virtual karaoke machine, a self-serve advertising platform, profile-editing tools, security systems, privacy filters, Myspace book lists, and on

and on." So instead of making its site an ever more comfortable and instinctive choice, Myspace kept its users off balance, wondering (if not subconsciously worrying) what was coming next.

Compare that with Facebook. From day one, Facebook has been building cumulative advantage. Initially it had some attractive features that Myspace lacked, making it a good value proposition, but more important to its success has been the consistency of its look and feel. Users conform to its rigid standards, and Facebook conforms to nothing or no one else. When it made its now-famous extension from desktop to mobile, the company ensured that users' mobile experience was highly consistent with their desktop experience.

To be sure, Facebook has from time to time introduced design changes in order to better leverage its functionality, and it has endured severe criticism in consequence. But in the main, new service introductions don't jeopardize comfort and familiarity, and the company has often made the changes optional in their initial stages. Even its name conjures up a familiar artifact, the college facebook, whereas Myspace gives the user no familiar reference at all.

Bottom line: By building on familiarity, Facebook has used cumulative advantage to become the most addictive social networking site in the world. That makes its subsidiary Instagram's decision to change its icon all the more baffling.

The Cumulative Advantage Imperatives

Myspace and Facebook nicely illustrate the twin realities that sustainable advantage is both possible and not assured. How, then, might the next Myspace enhance and extend its competitive edge by building a protective layer of cumulative advantage? Here are four basic rules to follow:

1. Become popular early

This idea is far from new—it is implicit in many of the best and earliest works on strategy, and we can see it in the thinking of Bruce Henderson, the founder of Boston Consulting Group. Henderson's particular focus was on the beneficial impact of cumulative output on costs—the now-famous experience curve, which suggests that as a company's experience in making something increases, its cost management becomes more efficient. He argued that companies should price aggressively early on—"ahead of the experience curve," in his parlance—and thus win sufficient market share to give the company lower costs, higher relative share, and higher profitability. The implication was clear: Early share advantage matters—a lot.

Marketers have long understood the importance of winning early. Launched specifically to serve the fast-growing automatic washing machine market, Tide is one of P&G's most revered, successful, and profitable brands. When it was introduced, in 1946, it immediately had the heaviest

advertising weight in the category. P&G also made sure that no washing machine was sold in America without a free box of Tide to get consumers' habits started. Tide quickly won the early popularity contest and has never looked back.

Free new-product samples to gain trial have always been a popular tactic with marketers. Aggressive pricing, the tactic favored by Henderson, is similarly popular. Samsung has emerged as the market share leader in the smartphone industry worldwide by providing very affordable Android-based phones that carriers can offer free with service contracts. For internet businesses, free is the core tactic for establishing habits. Virtually all the large-scale internet success stories—eBay, Google, Twitter, Instagram, Uber, Airbnb—make their services free so that users will grow and deepen their habits; then providers or advertisers will be willing to pay for access to them.

2. Design for habit

As we've seen, the best outcome is when choosing your offering becomes an automatic consumer response. So design for that—don't leave the outcome entirely to chance. We've seen how Facebook profits from its attention to consistent, habit-forming design, which has made use of its platform go beyond what we think of as habit: Checking for updates has become a real compulsion for a billion people. Of course Facebook benefits from increasingly huge network effects. But

the real advantage is that to switch from Facebook also entails breaking a powerful addiction.

The smartphone pioneer BlackBerry is perhaps the best example of a company that consciously designed for addiction. Its founder, Mike Lazaridis, explicitly created the device to make the cycle of feeling a buzz in the holster, slipping out the BlackBerry, checking the message, and thumbing a response on the miniature keyboard as addictive as possible. He succeeded: The device earned the nickname CrackBerry. The habit was so strong that even after BlackBerry had been brought down by the move to app-based and touch-screen smartphones, a core group of BlackBerry customers—who had staunchly refused to adapt—successfully implored the company's management to bring back a BlackBerry that resembled their previous-generation devices. It was given the comforting name Classic.

As Art Markman, a psychologist at the University of Texas, has pointed out to us, certain rules should be respected in designing for habit. To begin with, you must keep consistent those elements of the product design that can be seen from a distance so that buyers can find your product quickly. Distinctive colors and shapes like Tide's bright orange and the Doritos logo accomplish this.

And you should find ways to make products fit in people's environments to encourage use. When P&G introduced Febreze, consumers liked the way it worked but did not use it often. Part of the problem, it turned out, was that the container

was shaped like a glass-cleaner bottle, signaling that it should be kept under the sink. The bottle was ultimately redesigned to be kept on a counter or in a more visible cabinet, and use after purchase increased.

Unfortunately, the design changes that companies make all too often end up disrupting habits rather than strengthening them. Look for changes that will reinforce habits and encourage repurchase. The Amazon Dash Button provides an excellent example: By creating a simple way for people to reorder products they use often, Amazon helps them develop habits and locks them into a particular distribution channel.

3. Innovate inside the brand

As we've already noted, companies engage in initiatives to "relaunch," "repackage," or "replatform" at some peril: Such efforts can require customers to break their habits. Of course companies have to keep their products up-to-date, but changes in technology or other features should ideally be introduced in a manner that allows the new version of a product or service to retain the cumulative advantage of the old.

Even the most successful builders of cumulative advantage sometimes forget this rule. P&G, for example, which has increased Tide's cumulative advantage over 70 years through huge changes, has had to learn some painful lessons along the way. Arguably the first great detergent innovation after Tide's launch was the development of liquid detergents. P&G's first

response was to launch a new brand, called Era, in 1975. With no cumulative advantage behind it, Era failed to become a major brand despite consumers' increasing substitution of liquid for powdered detergent.

Recognizing that as the number one brand in the category, Tide had a strong connection with consumers and a powerful cumulative advantage, P&G decided to launch Liquid Tide in 1984, in familiar packaging and with consistent branding. It went on to become the dominant liquid detergent despite its late entry. After that experience, P&G was careful to ensure that further innovations were consistent with the Tide brand. When its scientists figured out how to incorporate bleach into detergent, the product was called Tide Plus Bleach. The breakthrough cold-cleaning technology appeared in Tide Coldwater, and the revolutionary three-in-one pod form was launched as Tide Pods. The branding could not have been simpler or clearer: This is your beloved Tide, with bleach added, for cold water, in pod form. These comfort- and familiarity-laden innovations reinforced rather than diminished the brand's cumulative advantage. The new products all preserved the look of Tide's traditional packaging—the brilliant orange and the bull's-eye logo. The few times in Tide history when that look was altered—such as with blue packaging for the Tide Coldwater launch—the effect on consumers was significantly negative, and the change was quickly reversed.

Of course, sometimes change is absolutely necessary to maintain relevance and advantage. In such situations smart companies succeed by helping customers transition from the old habit to the new one. Netflix began as a service that delivered DVDs to customers by mail. It would be out of business today if it had attempted to maximize continuity by refusing to change. Instead, it has successfully transformed itself into a video streaming service.

Although the new Netflix markets a completely different platform for digital entertainment, involving a new set of activities, Netflix found ways to help its customers by accentuating what did not have to change. It has the same look and feel and is still a subscription service that gives people access to the latest entertainment without leaving their homes. Thus its customers can deal with the necessary aspects of change while maintaining as much of the habit as possible. For customers, "improved" is much more comfortable and less scary than "new," however awesome "new" sounds to brand managers and advertising agencies.

4. Keep communication simple

One of the fathers of behavioral science, Daniel Kahneman, characterized subconscious, habit-driven decision making as "thinking fast" and conscious decision making as "thinking slow." Marketers and advertisers often seem to live in thinking-slow mode. They are rewarded with industry kudos for the

cleverness with which they weave together and highlight the multiple benefits of a new product or service. True, ads that are clever and memorable sometimes move customers to change their habits. The slow-thinking conscious mind, if it decides to pay attention, may well say, “Wow, that is impressive. I can’t wait!”

But if viewers aren’t paying attention (as in the vast majority of cases), an artful communication may backfire. Consider the ad that came out a couple of years ago for the Samsung Galaxy S5. It began by showing successive vignettes of generic-looking smartphones failing to (a) demonstrate water resistance; (b) protect against a young child’s accidentally sending an embarrassing message; and (c) enable an easy change of battery. It then triumphantly pointed out that the Samsung S5, which looked pretty much like the three previous phones, overcame all these flaws. Conscious, slow-thinking viewers, if they watched the whole ad, may have been persuaded that the S5 was different from and superior to other phones. But an arguably greater likelihood was that fast-thinking viewers would subconsciously associate the S5 with the three shortcomings. When making a purchase decision, they might be swayed by a subconscious plea: “Don’t buy the one with the water-resistance, rogue-message, and battery-change problems.” In fact, the ad might even induce them to buy a competitor’s product—such as the iPhone 7—whose message about water resistance is simpler to take in.

Remember: The mind is lazy. It doesn't want to ramp up attention to absorb a message with a high level of complexity. Simply showing the water resistance of the Samsung S5—or better yet, showing a customer buying an S5 and being told by the sales rep that it was fully water-resistant—would have been much more powerful. The latter would tell fast thinkers what you wanted them to do: go to a store and buy the Samsung S5. Of course, neither of those ads would be likely to win any awards from marketers focused on the cleverness of advertising copy.

Competitive Advantage Must Reads

EXPERTS HAVE BEEN DEBATING THE NATURE of competitive advantage for years. Below are four standout articles that articulate the most influential thinking on the subject. They can be found at HBR.org.

"What Is Strategy?" by Michael E. Porter. In this classic 1996 article, Porter argues that operational effectiveness, although necessary to superior performance, is not sufficient, because its techniques are easy to imitate. The essence of strategy is choosing a unique and valuable position rooted in activities that are much more difficult to match.

"The One Number You Need to Grow" by Frederick F. Reichheld. This 2003 article introduced the Net Promoter Score—a simple measure of a customer's willingness to recommend a

product. NPS is a reliable index to loyalty, says Reichheld, and the best predictor of top-line growth.

"Transient Advantage" by Rita Gunther McGrath. McGrath contends that business leaders are overly fixated on creating a sustainable competitive advantage. Business today is too turbulent to spend months crafting a long-term strategy, she says in this 2013 article. Rather, leaders need a portfolio of transient advantages that can be built quickly and abandoned just as rapidly.

"When Marketing *Is* Strategy" by Niraj Dawar. For decades, businesses have sought competitive advantage in upstream activities related to making new products—bigger factories, cheaper raw materials, efficiency, and so on. But those are all easily copied. Advantage, says Dawar in this 2013 article, increasingly lies in the marketplace. The important question is not "What else can we make?" but "What else can we do for our customers?"

The death of sustainable competitive advantage has been greatly exaggerated. Competitive advantage is as sustainable as it has always been. What is different today is that in a world of infinite communication and innovation, many strategists seem convinced that sustainability can be delivered only by constantly making a company's value proposition the conscious consumer's rational or emotional first choice. They have forgotten, or they never understood, the dominance of the subconscious mind in decision making. For fast thinkers,

products and services that are easy to access and that reinforce comfortable buying habits will over time trump innovative but unfamiliar alternatives that may be harder to find and require forming new habits.

So beware of falling into the trap of constantly updating your value proposition and branding. And any company, whether it is a large established player, a niche player, or a new entrant, can sustain the initial advantage provided by a superior value proposition by understanding and following the four rules of cumulative advantage.

Counterpoint

Old Habits Die Hard, but They Do Die

by Rita Gunther McGrath

I love the notion that customers' purchase decisions are more closely related to habit and ease than to loyalty—it brings much-needed insight from behavioral science to the study of consumer decisions. And, as Lafley and Martin suggest, it has major implications for how products are developed and brands are managed. I completely agree with the authors that

customers' unconscious minds dominate their decision-making process—and I suspect that any company can benefit from making their routine choices easier, faster, and more convenient. That's one reason the subscription model has become so popular in so many industries—it eliminates the need for customers to consciously decide about routine purchases and offers providers the lure of effortlessly recurring revenue.

The theory of cumulative advantage makes a lot of sense in what Martin Reeves and his colleagues at BCG call a *classical* strategic setting—one in which industry boundaries are clearly delineated, the basis of competition is stable, the environment experiences no major disruptions, and a strong competitive position, once created, can be sustained. As BCG has shown, the candy company Mars has enjoyed very long product life cycles: Snickers and M&M's (introduced in 1930 and 1941, respectively) are among the best-selling candies in the world today. Procter & Gamble has a similarly strong track record with Tide, Unilever with Dove, and PepsiCo with Tropicana orange juice.

But for a growing number of companies, those conditions don't apply. Their industry boundaries aren't clearly delineated—in fact, they're totally blurry. Just ask anyone in retail, entertainment, or telecommunications. Their environments aren't stable—companies can be disrupted by entrants from below, as Clayton Christensen has pointed out, but also by

competitors using a different business model or moving over from an adjacent industry. And long-standing competitive strengths can be upended almost overnight by someone who has digitized your physical business (hello, Encyclopaedia Britannica) or turned your product into a service (see Zipcar, Airbnb, and Uber). Apple and Google didn't necessarily *intend* to disrupt point-and-shoot cameras, stand-alone GPS devices, TV advertising, or the Weather Channel, but they did so nonetheless. (See the sidebar "It Works Until It Doesn't: The Changing Nature of Competitive Advantage.")

Strategic Inflection Points

For some time my argument has been that we need a new way of thinking about strategy in environments where traditional barriers to entry are eroding, or in which emerging technologies weaken constraints. Andy Grove's phrase *inflection point* captures this situation nicely. A strategic inflection point, he says, is "a time in the life of a business when its fundamentals are about to change." Inflection points are difficult for traditional strategy tools to address, because they usually don't look important at first. The Wright brothers proved it was possible to fly safely in 1903. Nobody took that seriously until 1908. Even with the 1914 launch of the first commercial flight, few realized that airplanes would upend industries as varied as railroads, steamships, and package delivery.

It Works Until It Doesn't: The Changing Nature of Competitive Advantage

ANY THEORY THAT SEEKS TO explain cause-and-effect relationships operates within a set of constraints. A theory that works beautifully under one set may fall apart under another.

Over the years, we have seen systematic shifts in how companies create a strategically valuable position, often reinforced by the constraints of the systems within which they operate. In the early 1900s, for instance, companies that achieved economies of scope and scale through mass production were dominant, and they remained so right through the period after World War II. Indeed, the *Fortune* 500 list of 1970 reveals the dominance of huge U.S.-based industrial players such as General Motors, General Electric, Exxon Mobil, and Union Carbide.

With the advent of communications and computational technology, strategic advantage began to shift toward companies that leveraged information technology to provide services in addition to goods, and toward models that placed a value on information utilization in addition to product features and functions. Although the industrial giants remained in place for a long time, companies such as Walmart, AIG, Enron, and Citigroup had joined them on the *Fortune* 500 list by 1995.

Today the dynamics of competitive advantage have shifted once more. Companies are achieving advantage through *access* to assets rather than ownership of them. In addition, a

whole new category of “platform” companies, such as Google, Apple, and Facebook, have emerged, and the very size of their customer base creates a reinforcing virtuous cycle. Often called network effects, these dynamics mean that the more customers a company has, the more valuable it is to each additional customer. In such cases being an early mover can result in a formidable advantage.

The point is that every theory has its constraints. Attempting to apply it outside those conditions can lead to disaster.

Consumer habits can be powerful aids to sustaining a competitive advantage, as Lafley and Martin quite correctly point out. But habits, like other elements of the environment, can change. And when new technologies make new business models viable, habits can change very fast.

Consider the powerful forces that were unleashed from 2004 to 2007 by four separate but linked business developments. In 2004 Facebook was founded. In 2005 YouTube was founded. In 2006 Amazon launched Amazon Web Services (AWS). In 2007 Apple’s iPhone and Google’s Android operating system were commercially released. As the technology analyst Ben Thompson points out, AWS made it easy and cheap to start an online company, YouTube made it easy and cheap to upload videos, and Facebook offered a ready-made channel for sharing such videos. I’d add that the wild popularity of mobile phones made all that available to ordinary people. Now a couple of guys with an idea and access to programming skills can rival

global giants in days or weeks, not months or years—with practically no assets.

Gillette Versus Dollar Shave

And that's exactly what happened with the 2012 launch of DollarShaveClub.com. The brand promise was simple: great razors with few frills, for a low subscription price, delivered to your door automatically. Not only did you save money, but you didn't have to visit a store or risk running out. This was all the more attractive because habitual buying behavior had already been disrupted: Razor blades are expensive and easy to steal, so it has become common for them to be kept under lock and key in stores. Today, although Dollar Shave Club has an 8% share of the $3 billion U.S. market for blades and razors, the far more important number is its "share of cartridge." That, according to recent sources, is an astonishing 15% of all cartridges sold.

In 2010 Gillette had 70% of the global shaving market and legions of loyal customers who reliably traded up as the next generation of products, with higher prices, were released. Procter & Gamble had acquired the brand in 2005 for a reported $57 billion. It was a classic high-market-share, high-quality business—and we can only assume from their track records that both Gillette and P&G were extremely good at getting customers to buy habitually. Clearly they had a strong

cumulative advantage. But that wasn't enough, because the business had hit an inflection point.

In July 2016 Unilever agreed to buy Dollar Shave Club for about $1 billion in cash. The founding entrepreneurs are happy. Their investors are happy. Their customers are clearly happy. The incumbents? Not so much. According to the *Wall Street Journal,* P&G's share of men's razors and blades had fallen to 59% in 2015. One of its responses was to launch the Gillette Shave Club. Having seen the potentially habit-destroying effects of the subscription model, P&G now offers subscription and delivery for other products—including expensive Tide Pods.

Twenty years ago it would have been inconceivable that a marketing message could reach 20 million people in a matter of weeks without massive spending on television and other advertising. But Dollar Shave Club accomplished that with an entertaining launch video, promotion on social media channels, and a group of enthusiastic brand ambassadors who provided feet on the ground to promote its products—free.

Leveraging the Familiar Even as You Reinvent

The point of this story is that even a company as storied as P&G can be taken by surprise. Which brings me to the tricky question, How can executives balance the formidable power of

cumulative advantage and habit, often associated with a brand, with the need to refresh their approach?

One practical tactic is to leverage the core skills or capabilities of an organization in a new format. Target offers an illustrative case. The company's roots were in a traditional department store, Dayton's, which became Dayton Hudson and eventually Marshall Field's. In 1960 its leadership saw an opportunity to reach a market segment that appeared to be growing but wasn't well served by the existing format. That segment consisted of value-conscious consumers who nonetheless appreciated good design and a reasonably pleasant shopping experience. To protect the then-dominant department store brand, the new venture was branded separately. Its iconic bull's-eye logo was meant to represent the notion of hitting the target of convenience, price, and customer experience.

By the mid-1970s Target stores were outselling the company's department stores. In 2000 Dayton Hudson changed its name to Target to reflect the reality of its now-core business. In 2004 the company sold its department store brands, completing an extraordinary retail transformation.

Another fascinating transformation that leveraged the core skills of a parent company is the relentless digitization pursued by the newspaper publisher Schibsted, of Norway. Unlike many other newspaper publishers, Schibsted saw the encroachment of digital classified advertisements as an opportunity rather than a threat to its business. Beginning in the late 1990s, its

leaders aggressively courted classified advertisers to list with its digital properties. This became a crusade. As Sverre Munck observed when he was the EVP for strategy and international editorial, “The Internet was made for classifieds and classifieds were made for the Internet.” Long a traditional media company, Schibsted was able to leverage deep ties with its advertisers with a model that permitted economies of scale in editorial and communication activities across its media brands. These were supplemented by a significant commitment to bringing technological capabilities into the very core of the media business, ending the tug-of-war between conventional editorial processes and the logic of digital transformation.

A Balance of Stability and Dynamism

In 2012 I wrote an HBR piece titled “How the Growth Outliers Do It.” That analysis, which looked at 10 years of net income data from 2000 to 2009, found that out of 2,347 of the publicly traded firms with a market capitalization of more than $1 billion, only 10 had successfully grown net income by 5% or more in every one of those 10 years. (Although performance can be measured in many ways, this seems to me to be one that tests the idea of sustainable advantage consistently.) The first conclusion is obvious: Steady, sustained profit growth is hard to achieve, particularly in a period that includes the Great Recession of 2008. The second, however, is that some

companies do manage to achieve it for relatively long periods of time. I found that those companies balanced elements of stability (culture, relationships, leadership, and even strategy) with elements of dynamism (rapid resource mobilization, marketplace experiments, and people mobility).

I spoke recently with Malcolm Frank, a senior executive at Cognizant, which appears on both my original list and one that I've updated through the end of 2015 (for which I used modified criteria: If a company was over the threshold for any year in the previous 10 years, it was included on the list, which totaled roughly 5,300). Frank told me that his organization lives and breathes the idea that in many cases competitive advantage is not going to last. "For us, what was the ceiling five years ago is going to be the floor five years from now," he said. Cognizant is also disciplined about exiting slow-growth or underperforming operations. But it is remarkably stable. Francisco D'Souza has been CEO since 2007, and the most recent addition to the leadership team joined in 2005. Cognizant's culture, too, reflects what its leaders call a "well-established set of cultural values," as demonstrated in their written documents, public statements, and go-to-market strategies.

But let's return to the really important insight that underlies the argument of Lafley and Martin: Most of the time, we are all

unaware of the true motivations behind the choices we make. The better strategists and marketers become at understanding those motivations, the more likely they are to succeed at building habitual behavior among consumers—and, just as important, the more likely they are to see how those habits might change. Clayton Christensen's "jobs to be done" theory may come in handy here. He has famously said that when we buy products, we are actually hiring them to do a job for us. And the "jobs" underlying most product purchases are remarkably stable. Take communication: From smoke signals to the Pony Express to the telegraph to the telephone to the communications technologies of today, our basic job—to send messages to other human beings—has not changed. But how that job gets done has changed dramatically. If incumbent companies stay focused on the job itself—rather than on the specifics of how it gets done at this moment in time—they may be able to invent a better way before the competition does.

This is a point that company leaders often miss. Customers can easily "hire" another solution that does a given job better—just as vast numbers of them are currently doing with razors bought by subscription.

Originally published in January–February 2017. Reprint R1701B

Noise

How to Overcome the High, Hidden Cost of Inconsistent Decision Making. *by Daniel Kahneman, Andrew M. Rosenfield, Linnea Gandhi, and Tom Blaser*

AT A GLOBAL FINANCIAL SERVICES FIRM we worked with, a longtime customer accidentally submitted the same application file to two offices. Though the employees who reviewed the file were supposed to follow the same guidelines—and thus arrive at similar outcomes—the separate offices returned very different quotes. Taken aback, the customer gave the business to a competitor. From the point of view of the firm, employees in the same role should have been interchangeable, but in this case they were not. Unfortunately, this is a common problem.

Professionals in many organizations are assigned arbitrarily to cases: appraisers in credit-rating agencies, physicians in emergency rooms, underwriters of loans and insurance, and others. Organizations expect consistency from these professionals: Identical cases should be treated similarly, if not identically. The problem is that humans are unreliable decision

makers; their judgments are strongly influenced by irrelevant factors, such as their current mood, the time since their last meal, and the weather. We call the chance variability of judgments *noise.* It is an invisible tax on the bottom line of many companies.

Some jobs are noise-free. Clerks at a bank or a post office perform complex tasks, but they must follow strict rules that limit subjective judgment and guarantee, by design, that identical cases will be treated identically. In contrast, medical professionals, loan officers, project managers, judges, and executives all make judgment calls, which are guided by informal experience and general principles rather than by rigid rules. And if they don't reach precisely the same answer that every other person in their role would, that's acceptable; this is what we mean when we say that a decision is "a matter of judgment." A firm whose employees exercise judgment does not expect decisions to be entirely free of noise. But often noise is *far above* the level that executives would consider tolerable—and they are completely unaware of it.

The prevalence of noise has been demonstrated in several studies. Academic researchers have repeatedly confirmed that professionals often contradict their own prior judgments when given the same data on different occasions. For instance, when software developers were asked on two separate days to estimate the completion time for a given task, the hours they projected differed by 71%, on average. When pathologists made

two assessments of the severity of biopsy results, the correlation between their ratings was only .61 (out of a perfect 1.0), indicating that they made inconsistent diagnoses quite frequently. Judgments made by different people are even more likely to diverge. Research has confirmed that in many tasks, experts' decisions are highly variable: valuing stocks, appraising real estate, sentencing criminals, evaluating job performance, auditing financial statements, and more. The unavoidable conclusion is that professionals often make decisions that deviate significantly from those of their peers, from their own prior decisions, and from rules that they themselves claim to follow.

Noise is often insidious: It causes even successful companies to lose substantial amounts of money without realizing it. How substantial? To get an estimate, we asked executives in one of the organizations we studied the following: "Suppose the optimal assessment of a case is $100,000. What would be the cost to the organization if the professional in charge of the case assessed a value of $115,000? What would be the cost of assessing it at $85,000?" The cost estimates were high. Aggregated over the assessments made every year, the cost of noise was measured in billions—an unacceptable number even for a large global firm. The value of reducing noise even by a few percentage points would be in the tens of millions. Remarkably, the organization had completely ignored the question of consistency until then.

Idea in Brief

The Problem

Many organizations expect consistency from their professional employees. However, human judgment is often influenced by such irrelevant factors as the weather and the last case seen. More important, decisions often vary from employee to employee. The chance variability of judgments is called *noise*, and it is surprisingly costly to companies.

The Starting Point

Managers should perform a noise audit in which members of a unit, working independently, evaluate a common set of cases. The degree to which their decisions vary is the measure of noise. It will often be dramatically higher than executives anticipate.

The Solution

The most radical solution to a severe noise problem is to replace human judgment with algorithms. Algorithms are not difficult to construct—but often they're politically or operationally infeasible. In such instances, companies should establish procedures to help professionals achieve greater consistency.

It has long been known that predictions and decisions generated by simple statistical algorithms are often more accurate than those made by experts, even when the experts have access to more information than the formulas use. It is

less well known that the key advantage of algorithms is that they are noise-free: Unlike humans, a formula will always return the same output for any given input. Superior consistency allows even simple and imperfect algorithms to achieve greater accuracy than human professionals. (Of course, there are times when algorithms will be operationally or politically infeasible, as we will discuss.)

In this article we explain the difference between noise and bias and look at how executives can audit the level and impact of noise in their organizations. We then describe an inexpensive, underused method for building algorithms that remediate noise, and we sketch out procedures that can promote consistency when algorithms are not an option.

Noise vs. Bias

When people consider errors in judgment and decision making, they most likely think of social biases like the stereotyping of minorities or of cognitive biases such as overconfidence and unfounded optimism. The useless variability that we call noise is a different type of error. To appreciate the distinction, think of your bathroom scale. We would say that the scale is *biased* if its readings are generally either too high or too low. If your weight appears to depend on where you happen to place your feet, the scale is *noisy.* A scale that consistently underestimates true weight by exactly four pounds is seriously biased but free

of noise. A scale that gives two different readings when you step on it twice is noisy. Many errors of measurement arise from a combination of bias and noise. Most inexpensive bathroom scales are somewhat biased and quite noisy.

For a visual illustration of the distinction, consider the targets in the exhibit "How noise and bias affect accuracy." These show the results of target practice for four-person teams in which each individual shoots once.

- Team A is *accurate:* The shots of the teammates are on the bull's-eye and close to one another.
- The other three teams are inaccurate but in distinctive ways:
- Team B is *noisy:* The shots of its members are centered around the bull's-eye but widely scattered.
- Team C is *biased:* The shots all missed the bull's-eye but cluster together.
- Team D is both *noisy* and *biased.*

As a comparison of teams A and B illustrates, an increase in noise always impairs accuracy when there is no bias. When bias is present, increasing noise may actually cause a lucky hit, as happened for team D. Of course, no organization would put its trust in luck. Noise is always undesirable—and sometimes disastrous.

How noise and bias affect accuracy

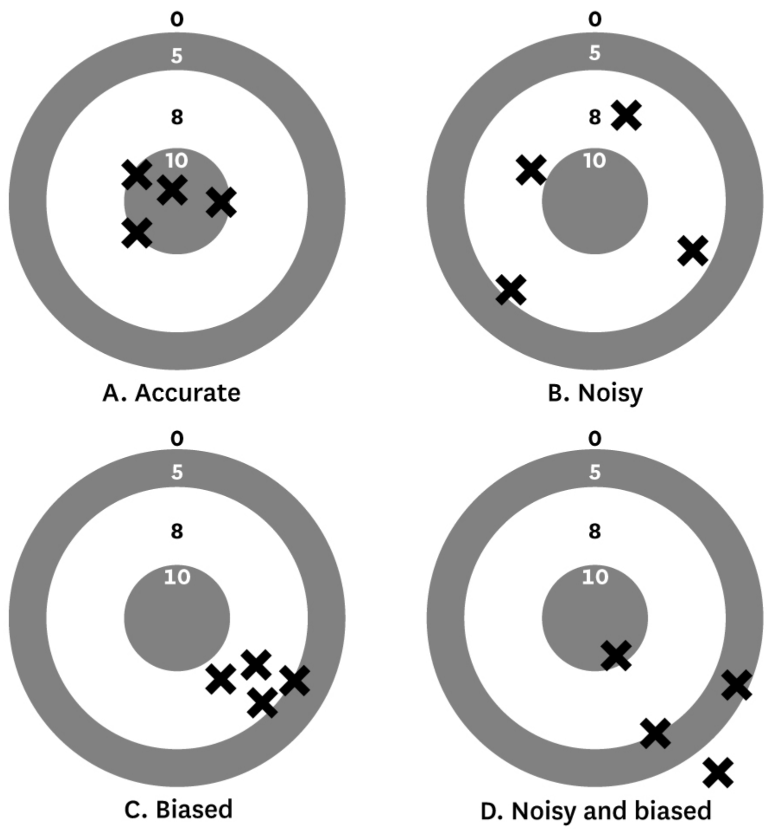

It is obviously useful to an organization to know about bias and noise in the decisions of its employees, but collecting that information isn't straightforward. Different issues arise in measuring these errors. A major problem is that the outcomes of decisions often aren't known until far in the future, if at all. Loan officers, for example, frequently must wait several years to see how loans they approved worked out, and they almost never know what happens to an applicant they reject.

Unlike bias, noise can be measured without knowing what an accurate response would be. To illustrate, imagine that the

targets at which the shooters aimed were erased from the exhibit. You would know nothing about the teams' overall accuracy, but you could be certain that something was wrong with the scattered shots of teams B and D: Wherever the bull's-eye was, they did not all come close to hitting it. All that's required to measure noise in judgments is a simple experiment in which a few realistic cases are evaluated independently by several professionals. Here again, the scattering of judgments can be observed without knowing the correct answer. We call such experiments *noise audits*.

Performing a Noise Audit

The point of a noise audit is not to produce a report. The ultimate goal is to improve the quality of decisions, and an audit can be successful only if the leaders of the unit are prepared to accept unpleasant results and act on them. Such buy-in is easier to achieve if the executives view the study as their own creation. To that end, the cases should be compiled by respected team members and should cover the range of problems typically encountered. To make the results relevant to everyone, all unit members should participate in the audit. A social scientist with experience in conducting rigorous behavioral experiments should supervise the technical aspects of the audit, but the professional unit must own the process.

Recently, we helped two financial services organizations conduct noise audits. The duties and expertise of the two groups we studied were quite different, but both required the evaluation of moderately complex materials and often involved decisions about hundreds of thousands of dollars. We followed the same protocol in both organizations. First we asked managers of the professional teams involved to construct several realistic case files for evaluation. To prevent information about the experiment from leaking, the entire exercise was conducted on the same day. Employees were asked to spend about half the day analyzing two to four cases. They were to decide on a dollar amount for each, as in their normal routine. To avoid collusion, the participants were not told that the study was concerned with reliability. In one organization, for example, the goals were described as understanding the employees' professional thinking, increasing their tools' usefulness, and improving communication among colleagues. About 70 professionals in organization A participated, and about 50 in organization B.

We constructed a noise index for each case, which answered the following question: "By how much do the judgments of two randomly chosen employees differ?" We expressed this amount as a percentage of their average. Suppose the assessments of a case by two employees are $600 and $1,000. The average of their assessments is $800, and the difference between them is $400, so the noise index is 50% for this pair.

We performed the same computation for all pairs of employees and then calculated an overall average noise index for each case.

Pre-audit interviews with executives in the two organizations indicated that they expected the differences between their professionals' decisions to range from 5% to 10%—a level they considered acceptable for "matters of judgment." The results came as a shock. The noise index ranged from 34% to 62% for the six cases in organization A, and the overall average was 48%. In the four cases in organization B, the noise index ranged from 46% to 70%, with an average of 60%. Perhaps most disappointing, experience on the job did not appear to reduce noise. Among professionals with five or more years on the job, average disagreement was 46% in organization A and 62% in organization B.

No one had seen this coming. But because they owned the study, the executives in both organizations accepted the conclusion that the judgments of their professionals were unreliable to an extent that could not be tolerated. All quickly agreed that something had to be done to control the problem.

Because the findings were consistent with prior research on the low reliability of professional judgment, they didn't surprise us. The major puzzle for us was the fact that neither organization had ever considered reliability to be an issue.

The problem of noise is effectively invisible in the business world; we have observed that audiences are quite surprised

when the reliability of professional judgment is mentioned as an issue. What prevents companies from recognizing that the judgments of their employees are noisy? The answer lies in two familiar phenomena: Experienced professionals tend to have high confidence in the accuracy of their own judgments, and they also have high regard for their colleagues' intelligence. This combination inevitably leads to an overestimation of agreement. When asked about what their colleagues would say, professionals expect others' judgments to be much closer to their own than they actually are. Most of the time, of course, experienced professionals are completely unconcerned with what others might think and simply assume that theirs is the best answer. One reason the problem of noise is invisible is that people do not go through life imagining plausible alternatives to every judgment they make.

The expectation that others will agree with you is sometimes justified, particularly where judgments are so skilled that they are intuitive. High-level chess and driving are standard examples of tasks that have been practiced to near perfection. Master players who look at a situation on a chessboard will all have very similar assessments of the state of the game—whether, say, the white queen is in danger or black's king-side defense is weak. The same is true of drivers. Negotiating traffic would be impossibly dangerous if we could not assume that the drivers around us share our understanding of priorities at

intersections and roundabouts. There is little or no noise at high levels of skill.

High skill develops in chess and driving through years of practice in a predictable environment, in which actions are followed by feedback that is both immediate and clear. Unfortunately, few professionals operate in such a world. In most jobs people learn to make judgments by hearing managers and colleagues explain and criticize—a much less reliable source of knowledge than learning from one's mistakes. Long experience on a job always increases people's confidence in their judgments, but in the absence of rapid feedback, confidence is no guarantee of either accuracy or consensus.

We offer this aphorism in summary: *Where there is judgment, there is noise—and usually more of it than you think*. As a rule, we believe that neither professionals nor their managers can make a good guess about the reliability of their judgments. The only way to get an accurate assessment is to conduct a noise audit. And at least in some cases, the problem will be severe enough to require action.

Types of noise and bias

Bias and noise are distinct kinds of error. Each comes in different variants and requires different corrective actions.

Type of bias	Examples	Corrective actions
General The average judgment is wrong.	• Planning fallacy: Forecasts of outcomes are mostly optimistic • Excessive risk aversion: A venture capital firm rejects too many promising but risky investments	• Continual monitoring of decisions • Guidelines and targets for the frequency of certain outcomes (such as loan approvals) • Eliminating incentives that favor biases

Type of bias	Examples	Corrective actions
Social Discrimination occurs against—or for—certain categories of cases.	• Frequent denial of credit to qualified applicants from certain ethnic groups • Gender bias in assessments of job performance	• Monitoring statistics for different groups • Blinding of applications • Objective and quantifiable metrics • Open channels for complaints • Guidelines and training

Type of bias	Examples	Corrective actions
Cognitive Decisions are strongly influenced by irrelevant factors or insensitive to relevant ones.	• Excessive effects of first impressions • Effects of anchors (such as an opening offer in negotiation) • Myopic neglect of future consequences	• Training employees to detect situations in which biases are likely to occur • Critiques of important decisions, focused on likely biases

Type of noise	Examples	Corrective actions
Variability across occasions Decisions vary when the same case is presented more than once to the same individual.	• A hiring officer's judgments of a file are influenced by her mood or the quality of the previous applicant	• Algorithms to replace human judgment • Checklists that encourage a consistent approach to decisions

Type of bias	Examples	Corrective actions
Variability across individuals Professionals in the same role make different decisions.	• Some individuals are generally more lenient than others • Some individuals are more cautious than others	• Algorithms to replace human judgment • Frequent monitoring of individuals' decisions • Roundtables at which differences are explored and resolved • Checklists that encourage a consistent approach to decisions

Dialing Down the Noise

The most radical solution to the noise problem is to replace human judgment with formal rules—known as algorithms—that use the data about a case to produce a prediction or a

decision. People have competed against algorithms in several hundred contests of accuracy over the past 60 years, in tasks ranging from predicting the life expectancy of cancer patients to predicting the success of graduate students. Algorithms were more accurate than human professionals in about half the studies, and approximately tied with the humans in the others. The ties should also count as victories for the algorithms, which are more cost-effective.

In many situations, of course, algorithms will not be practical. The application of a rule may not be feasible when inputs are idiosyncratic or hard to code in a consistent format. Algorithms are also less likely to be useful for judgments or decisions that involve multiple dimensions or depend on negotiation with another party. Even when an algorithmic solution is available in principle, organizational considerations sometimes prevent implementation. The replacement of existing employees by software is a painful process that will encounter resistance unless it frees those employees up for more-enjoyable tasks.

But if the conditions are right, developing and implementing algorithms can be surprisingly easy. The common assumption is that algorithms require statistical analysis of large amounts of data. For example, most people we talk to believe that data on thousands of loan applications and their outcomes is needed to develop an equation that predicts commercial loan defaults. Very few know that adequate algorithms can be developed

without any outcome data at all—and with input information on only a small number of cases. We call predictive formulas that are built without outcome data "reasoned rules," because they draw on commonsense reasoning.

The construction of a reasoned rule starts with the selection of a few (perhaps six to eight) variables that are incontrovertibly related to the outcome being predicted. If the outcome is loan default, for example, assets and liabilities will surely be included in the list. The next step is to assign these variables equal weight in the prediction formula, setting their sign in the obvious direction (positive for assets, negative for liabilities). The rule can then be constructed by a few simple calculations. (For more details, see the sidebar "How to Build a Reasoned Rule.")

How to Build a Reasoned Rule

YOU DON'T NEED OUTCOME DATA to create useful predictive algorithms. For example, you can build a reasoned rule that predicts loan defaults quite effectively without knowing what happened to past loans; all you need is a small set of recent loan applications. Here are the next steps:

1. Select six to eight variables that are distinct and obviously related to the predicted outcome. Assets and revenues (weighted positively) and liabilities (weighted negatively) would surely be included, along with a few other features of loan applications.

2. Take the data from your set of cases (all the loan applications from the past year) and compute the mean and standard deviation of each variable in that set.
3. For every case in the set, compute a “standard score” for each variable: the difference between the value in the case and the mean of the whole set, divided by the standard deviation. With standard scores, all variables are expressed on the same scale and can be compared and averaged.
4. Compute a “summary score” for each case—the average of its variables’ standard scores. This is the output of the reasoned rule. The same formula will be used for new cases, using the mean and standard deviation of the original set and updating periodically.
5. Order the cases in the set from high to low summary scores, and determine the appropriate actions for different ranges of scores. With loan applications, for instance, the actions might be “the top 10% of applicants will receive a discount” and “the bottom 30% will be turned down.”

You are now ready to apply the rule to new cases. The algorithm will compute a summary score for each new case and generate a decision.

The surprising result of much research is that in many contexts reasoned rules are about as accurate as statistical models built with outcome data. Standard statistical models combine a set of predictive variables, which are assigned

weights based on their relationship to the predicted outcomes and to one another. In many situations, however, these weights are both statistically unstable and practically unimportant. A simple rule that assigns equal weights to the selected variables is likely to be just as valid. Algorithms that weight variables equally and don't rely on outcome data have proved successful in personnel selection, election forecasting, predictions about football games, and other applications.

The bottom line here is that if you plan to use an algorithm to reduce noise, you need not wait for outcome data. You can reap most of the benefits by using common sense to select variables and the simplest possible rule to combine them.

Of course, no matter what type of algorithm is employed, people must retain ultimate control. Algorithms must be monitored and adjusted for occasional changes in the population of cases. Managers must also keep an eye on individual decisions and have the authority to override the algorithm in clear-cut cases. For example, a decision to approve a loan should be provisionally reversed if the firm discovers that the applicant has been arrested. Most important, executives should determine how to translate the algorithm's output into action. The algorithm can tell you which prospective loans are in the top 5% or in the bottom 10% of all applications, but someone must decide what to do with that information.

Algorithms are sometimes used as an intermediate source of information for professionals, who make the final decisions. One example is the Public Safety Assessment, a formula that was developed to help U.S. judges decide whether a defendant can be safely released pending trial. In its first six months of use in Kentucky, crime among defendants on pretrial release fell by about 15%, while the percentage of people released pretrial increased. It's obvious in this case that human judges must retain the final authority for the decisions: The public would be shocked to see justice meted out by a formula.

Uncomfortable as people may be with the idea, studies have shown that while humans can provide useful input to formulas, algorithms do better in the role of final decision maker. If the avoidance of errors is the only criterion, managers should be strongly advised to overrule the algorithm only in exceptional circumstances.

Bringing Discipline to Judgment

Replacing human decisions with an algorithm should be considered whenever professional judgments are noisy, but in most cases this solution will be too radical or simply impractical. An alternative is to adopt procedures that promote consistency by ensuring that employees in the same role use similar methods to seek information, integrate it into a view of the case, and translate that view into a decision. A thorough

examination of everything required to do that is beyond the scope of this article, but we can offer some basic advice, with the important caveat that instilling discipline in judgment is not at all easy.

Training is crucial, of course, but even professionals who were trained together tend to drift into their own way of doing things. Firms sometimes combat drift by organizing roundtables at which decision makers gather to review cases. Unfortunately, most roundtables are run in a way that makes it much too easy to achieve agreement, because participants quickly converge on the opinions stated first or most confidently. To prevent such spurious agreement, the individual participants in a roundtable should study the case independently, form opinions they're prepared to defend, and send those opinions to the group leader before the meeting. Such roundtables will effectively provide an audit of noise, with the added step of a group discussion in which differences of opinion are explored.

As an alternative or addition to roundtables, professionals should be offered user-friendly tools, such as checklists and carefully formulated questions, to guide them as they collect information about a case, make intermediate judgments, and formulate a final decision. Unwanted variability occurs at each of those stages, and firms can—and should—test how much such tools reduce it. Ideally, the people who use these tools will view them as aids that help them do their jobs effectively and

economically. Unfortunately, our experience suggests that the task of constructing judgment tools that are both effective and user-friendly is more difficult than many executives think. Controlling noise is hard, but we expect that an organization that conducts an audit and evaluates the cost of noise in dollars will conclude that reducing random variability is worth the effort.

Our main goal in this article is to introduce managers to the concept of noise as a source of errors and explain how it is distinct from bias. The term "bias" has entered the public consciousness to the extent that the words "error" and "bias" are often used interchangeably. In fact, better decisions are not achieved merely by reducing general biases (such as optimism) or specific social and cognitive biases (such as discrimination against women or anchoring effects). Executives who are concerned with accuracy should also confront the prevalence of inconsistency in professional judgments. Noise is more difficult to appreciate than bias, but it is no less real or less costly.

Originally published in October 2016. Reprint R1610B

Visualizations That Really Work

by Scott Berinato

NOT LONG AGO, THE ABILITY to create smart data visualizations, or dataviz, was a nice-to-have skill. For the most part, it benefited design- and data-minded managers who made a deliberate decision to invest in acquiring it. That's changed. Now visual communication is a must-have skill for all managers, because more and more often, it's the only way to make sense of the work they do.

Data is the primary force behind this shift. Decision making increasingly relies on data, which comes at us with such overwhelming velocity, and in such volume, that we can't comprehend it without some layer of abstraction, such as a visual one. A typical example: At Boeing the managers of the Osprey program need to improve the efficiency of the aircraft's takeoffs and landings. But each time the Osprey gets off the ground or touches back down, its sensors create a terabyte of data. Ten takeoffs and landings produce as much data as is held in the Library of Congress. Without visualization,

detecting the inefficiencies hidden in the patterns and anomalies of that data would be an impossible slog.

But even information that's not statistical demands visual expression. Complex systems—business process workflows, for example, or the way customers move through a store—are hard to understand, much less fix, if you can't first see them.

Thanks to the internet and a growing number of affordable tools, translating information into visuals is now easy (and cheap) for everyone, regardless of data skills or design skills. This is largely a positive development. One drawback, though, is that it reinforces the impulse to "click and viz" without first thinking about your purpose and goals. *Convenient* is a tempting replacement for good, but it will lead to charts that are merely adequate or, worse, ineffective. Automatically converting spreadsheet cells into a chart only visualizes pieces of a spreadsheet; it doesn't capture an idea. As the presentation expert Nancy Duarte puts it, "Don't project the idea that you're showing a chart. Project the idea that you're showing a reflection of human activity, of things people did to make a line go up and down. It's not 'Here are our Q3 financial results,' it's 'Here's where we missed our targets.'"

Managers who want to get better at making charts often start by learning rules. When should I use a bar chart? How many colors are too many? Where should the key go? Do I have to start my y-axis at zero? Visual grammar is important and useful —but knowing it doesn't guarantee that you'll make good

charts. To start with chart-making rules is to forgo strategy for execution; it's to pack for a trip without knowing where you're going.

Your visual communication will prove far more successful if you begin by acknowledging that it is not a lone action but, rather, several activities, each of which requires distinct types of planning, resources, and skills. The typology I offer here was created as a reaction to my making the very mistake I just described: The book from which this article is adapted started out as something like a rule book. But after exploring the history of visualization, the exciting state of visualization research, and smart ideas from experts and pioneers, I reconsidered the project. We didn't need another rule book; we needed a way to think about the increasingly crucial discipline of visual communication as a whole.

The typology described in this article is simple. By answering just two questions, you can set yourself up to succeed.

Idea in Brief

Context

Knowledge workers need greater visual literacy than they used to, because so much data—and so many ideas—are now presented graphically. But few of us have been taught data-visualization skills.

Tools are fine ...

Inexpensive tools allow anyone to perform simple tasks such as importing spreadsheet data into a bar chart. But that means it's easy to create terrible charts. Visualization can be so much more: It's an agile, powerful way to explore ideas and communicate information.

... But strategy is key

Don't jump straight to execution. Instead, first think about what you're representing—ideas or data? Then consider your purpose: Do you want to inform, persuade, or explore? The answers will suggest what tools and resources you need.

The Two Questions

To start thinking visually, consider the nature and purpose of your visualization:

Is the information *conceptual* or *data-driven?*

Am I *declaring* something or *exploring* something?

If you know the answers to these questions, you can plan what resources and tools you'll need and begin to discern what type of visualization will help you achieve your goals most effectively.

	Conceptual	Data-driven
Focus	*Ideas*	*Statistics*

	Conceptual	Data-driven
Goals	***Simplify, teach*** "Here's how our organization is structured."	***Inform, enlighten*** "Here are our revenues for the past two years."

The first question is the simpler of the two, and the answer is usually obvious. Either you're visualizing qualitative information or you're plotting quantitative information: ideas or statistics. But notice that the question is about the information itself, not the forms you might ultimately use to show it. For example, the classic Gartner Hype Cycle (see following page) uses a traditionally data-driven form—a line chart—but no actual data. It's a concept.

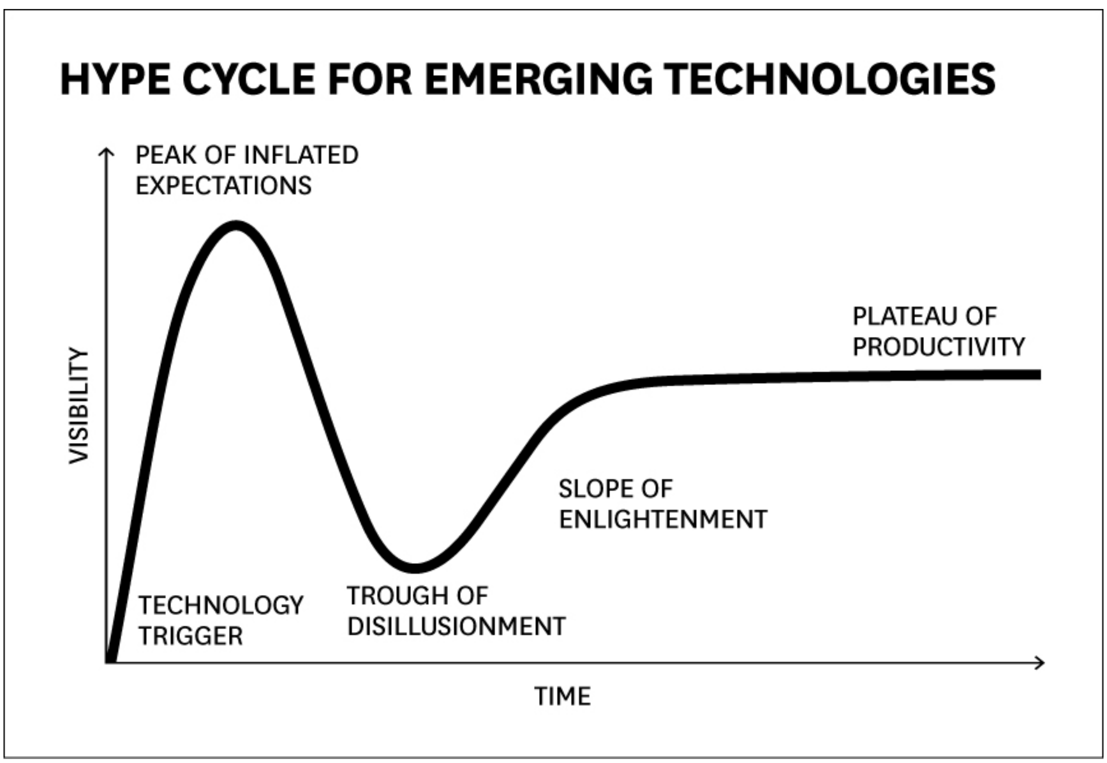

If the first question identifies what you *have*, the second elicits what you're *doing*: either communicating information (declarative) or trying to figure something out (exploratory).

	Declarative	Exploratory
Focus	*Documenting, designing*	*Prototyping, iterating, interacting, automating*
Goals	*Affirm* "Here is our budget by department."	*Confirm* "Let's see if marketing investments contributed to rising profits." *Discover* "What would we see if we visualized customer purchases by gender, location, and purchase amount in real time?"

Managers most often work with declarative visualizations, which make a statement, usually to an audience in a formal setting. If you have a spreadsheet workbook full of sales data and you're using it to show quarterly sales in a presentation, your purpose is declarative.

But let's say your boss wants to understand why the sales team's performance has lagged lately. You suspect that seasonal cycles have caused the dip, but you're not sure. Now your purpose is exploratory, and you'll use the same data to create visuals that will confirm or refute your hypothesis. The audience is usually yourself or a small team. If your hypothesis

is confirmed, you may well show your boss a declarative visualization, saying, “Here’s what’s happening to sales.”

Exploratory visualizations are actually of two kinds. In the example above, you were testing a hypothesis. But suppose you don’t have an idea about why performance is lagging—you don’t know what you’re looking for. You want to mine your workbook to see what patterns, trends, and anomalies emerge. What will you see, for example, when you measure sales performance in relation to the size of the region a salesperson manages? What happens if you compare seasonal trends in various geographies? How does weather affect sales? Such data brainstorming can deliver fresh insights. Big strategic questions—Why are revenues falling? Where can we find efficiencies? How do customers interact with us?—can benefit from a discovery-focused exploratory visualization.

The Four Types

The nature and purpose questions combine in a classic 2×2 to define four types of visual communication: idea illustration, idea generation, visual discovery, and everyday dataviz.

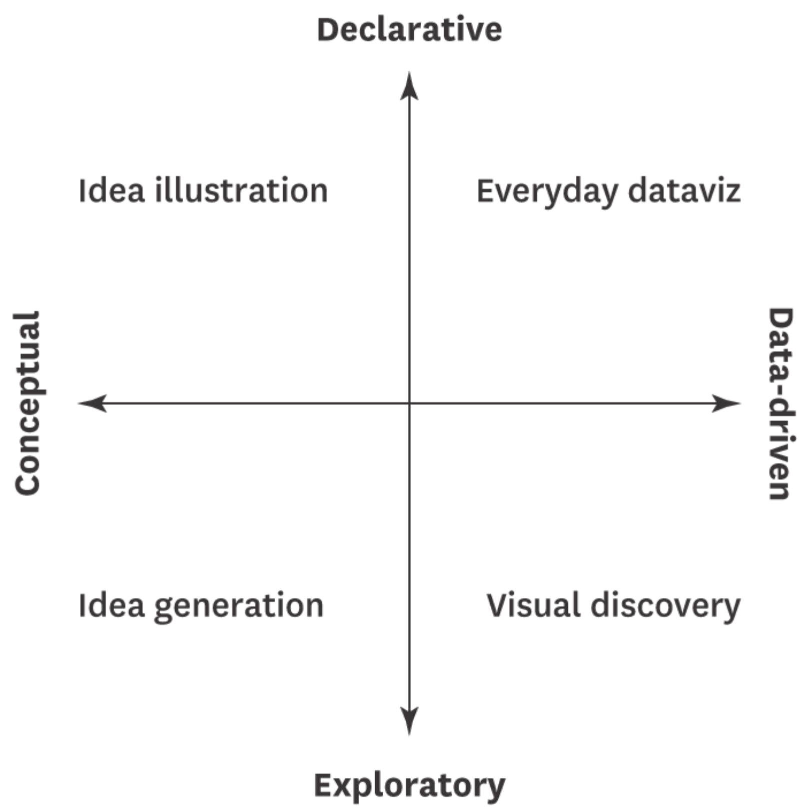

Idea illustration

Info type	Process, framework
Typical setting	Presentations, teaching
Primary skills	Design, editing
Goals	Learning, simplifying, explaining

We might call this quadrant the “consultants’ corner.” Consultants can’t resist process diagrams, cycle diagrams, and the like. At their best, idea illustrations clarify complex ideas by drawing on our ability to understand metaphors (trees, bridges) and simple design conventions (circles, hierarchies). Org charts

and decision trees are classic examples of idea illustration. So is the 2×2 that frames this article.

Idea illustration demands clear and simple design, but its reliance on metaphor invites unnecessary adornment. Because the discipline and boundaries of data sets aren't built in to idea illustration, they must be imposed. The focus should be on clear communication, structure, and the logic of the ideas. The most useful skills here are similar to what a text editor brings to a manuscript—the ability to pare things down to their essence. Some design skills will be useful too, whether they're your own or hired.

Suppose a company engages consultants to help its R&D group find inspiration in other industries. The consultants use a technique called the *pyramid search*—a way to get information from experts in other fields close to your own, who point you to the top experts in their fields, who point you to experts in still other fields, who then help you find the experts in those fields, and so on.

It's actually tricky to explain, so the consultants may use visualization to help. How does a pyramid search work? It looks something like this:

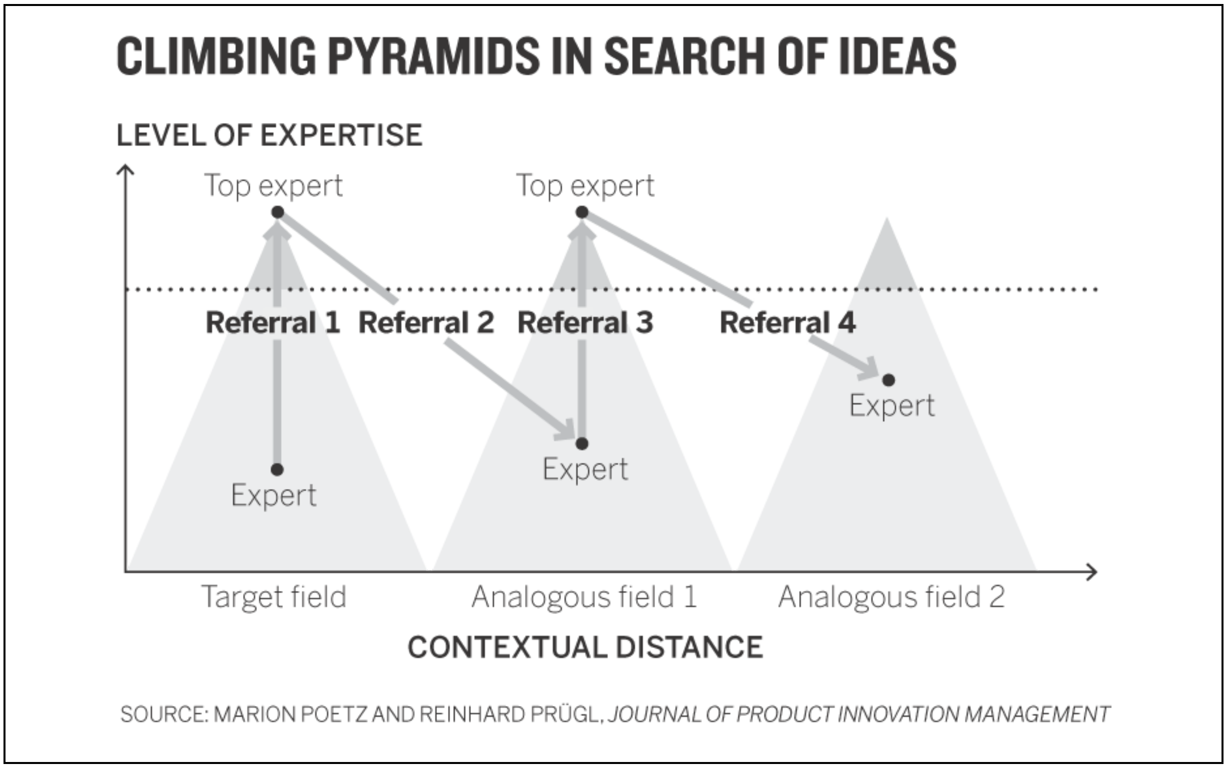

The axes use conventions that we can grasp immediately: industries plotted near to far and expertise mapped low to high. The pyramid shape itself shows the relative rarity of top experts compared with lower-level ones. Words in the title —“climbing” and “pyramids”—help us grasp the idea quickly. Finally, the designer didn’t succumb to a temptation to decorate: The pyramids aren’t literal, three-dimensional, sandstone-colored objects.

Too often, idea illustration doesn’t go that well, and you end up with something like this:

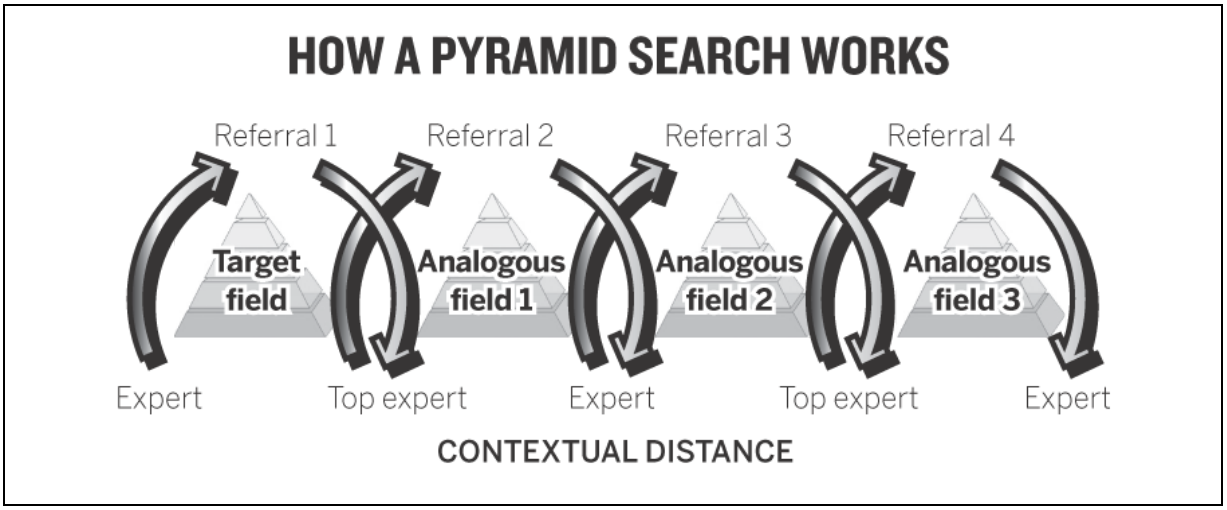

Here the color gradient, the drop shadows, and the 3-D pyramids distract us from the idea. The arrows don't actually demonstrate how a pyramid search works. And experts and top experts are placed on the same plane instead of at different heights to convey relative status.

Idea generation

Info type	Complex, undefined
Typical setting	Working session, brainstorming
Primary skills	Team-building, facilitation
Goals	Problem solving, discovery, innovation

Managers may not think of visualization as a tool to support idea generation, but they use it to brainstorm all the time—on whiteboards, on butcher paper, or, classically, on the back of a napkin. Like idea illustration, idea generation relies on conceptual metaphors, but it takes place in more-informal settings, such as off-sites, strategy sessions, and early-phase

innovation projects. It's used to find new ways of seeing how the business works and to answer complex managerial challenges: restructuring an organization, coming up with a new business process, codifying a system for making decisions.

Although idea generation can be done alone, it benefits from collaboration and borrows from design thinking—gathering as many diverse points of view and visual approaches as possible before homing in on one and refining it. Jon Kolko, the founder and director of the Austin Center for Design and the author of *Well-Designed: How to Use Empathy to Create Products People Love,* fills the whiteboard walls of his office with conceptual, exploratory visualizations. "It's our go-to method for thinking through complexity," he says. "Sketching is this effort to work through ambiguity and muddiness and come to crispness." Managers who are good at leading teams, facilitating brainstorming sessions, and encouraging and then capturing creative thinking will do well in this quadrant. Design skills and editing are less important here, and sometimes counterproductive. When you're seeking breakthroughs, editing is the opposite of what you need, and you should think in rapid sketches; refined designs will just slow you down.

Suppose a marketing team is holding an off-site. The team members need to come up with a way to show executives their proposed strategy for going upmarket. An hourlong whiteboard session yields several approaches and ideas (none of which are erased) for presenting the strategy. Ultimately, one approach

gains purchase with the team, which thinks it best captures the key point: Get fewer customers to spend much more. The whiteboard looks something like this:

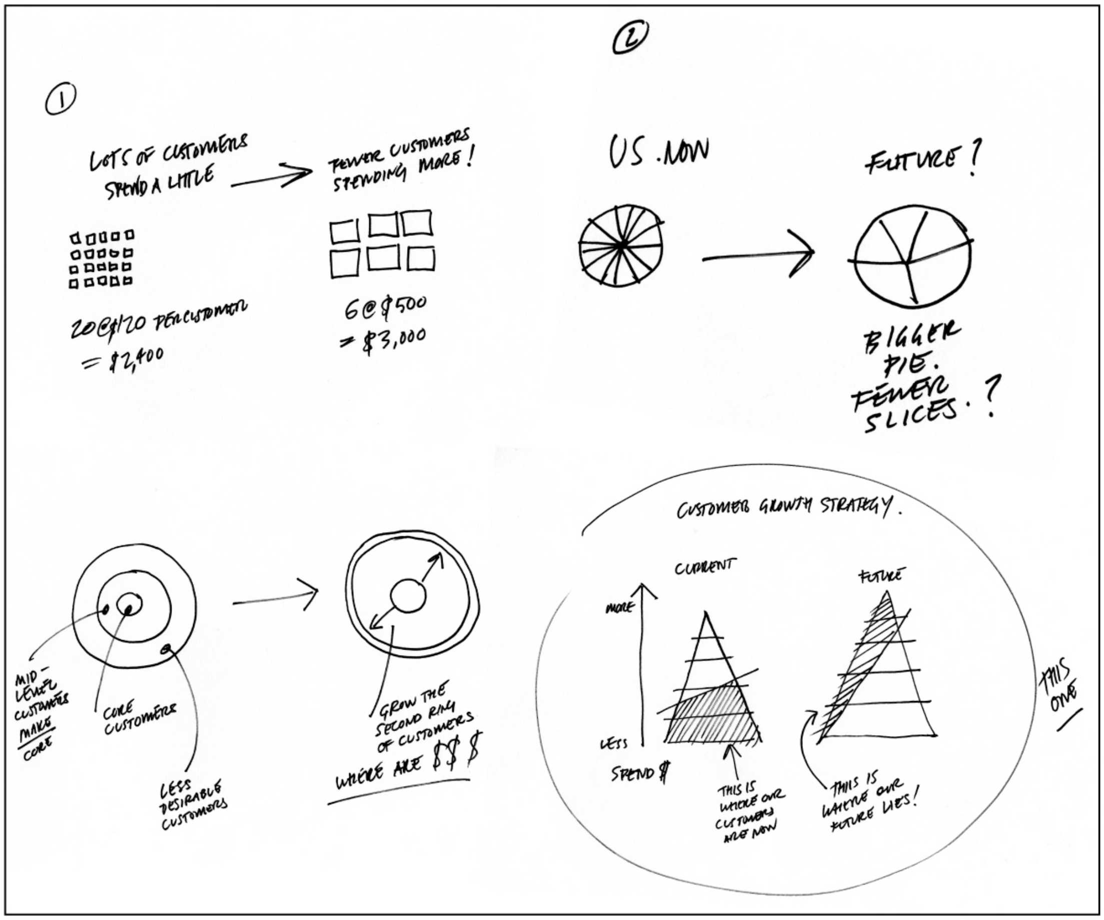

Of course, visuals that emerge from idea generation often lead to more formally designed and presented idea illustrations.

Visual discovery

Info type	Big data, complex, dynamic
Typical	Working sessions, testing, analysis

setting	
Primary skills	Business intelligence, programming, paired analysis
Goals	Trend spotting, sense making, deep analysis

This is the most complicated quadrant, because in truth it holds two categories. Recall that we originally separated exploratory purposes into two kinds: testing a hypothesis and mining for patterns, trends, and anomalies. The former is focused, whereas the latter is more flexible. The bigger and more complex the data, and the less you know going in, the more open-ended the work.

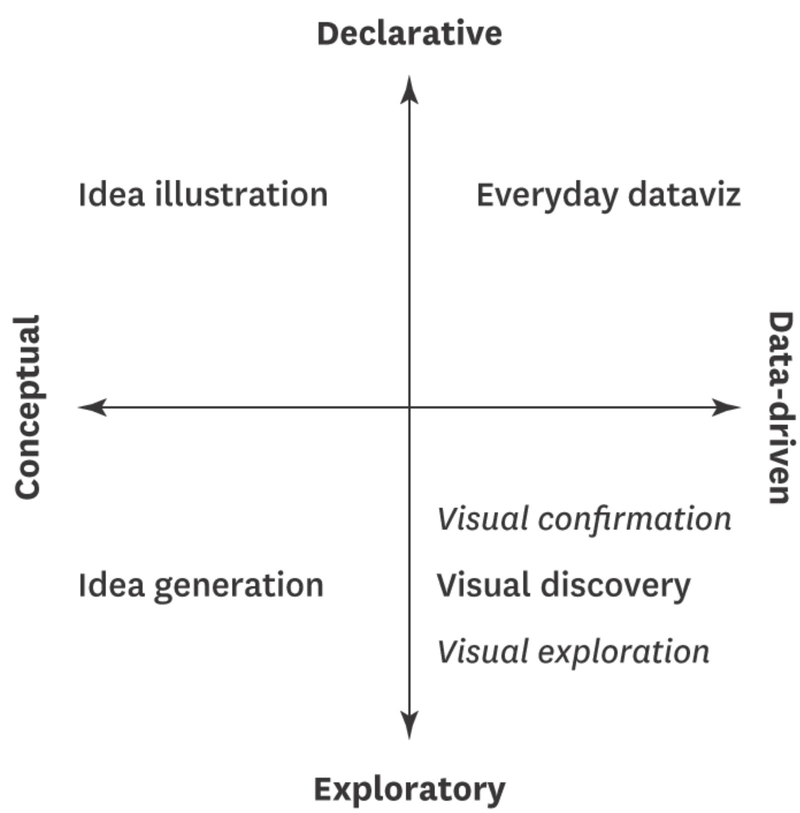

Visual Confirmation. You're answering one of two questions with this kind of project: Is what I suspect actually true? or

What are some other ways of depicting this idea?

The scope of the data tends to be manageable, and the chart types you're likely to use are common—although when trying to depict things in new ways, you may venture into some less-common types. Confirmation usually doesn't happen in a formal setting; it's the work you do to find the charts you want to create for presentations. That means your time will shift away from design and toward prototyping that allows you to rapidly iterate on the dataviz. Some skill at manipulating spreadsheets and knowledge of programs or sites that enable swift prototyping are useful here.

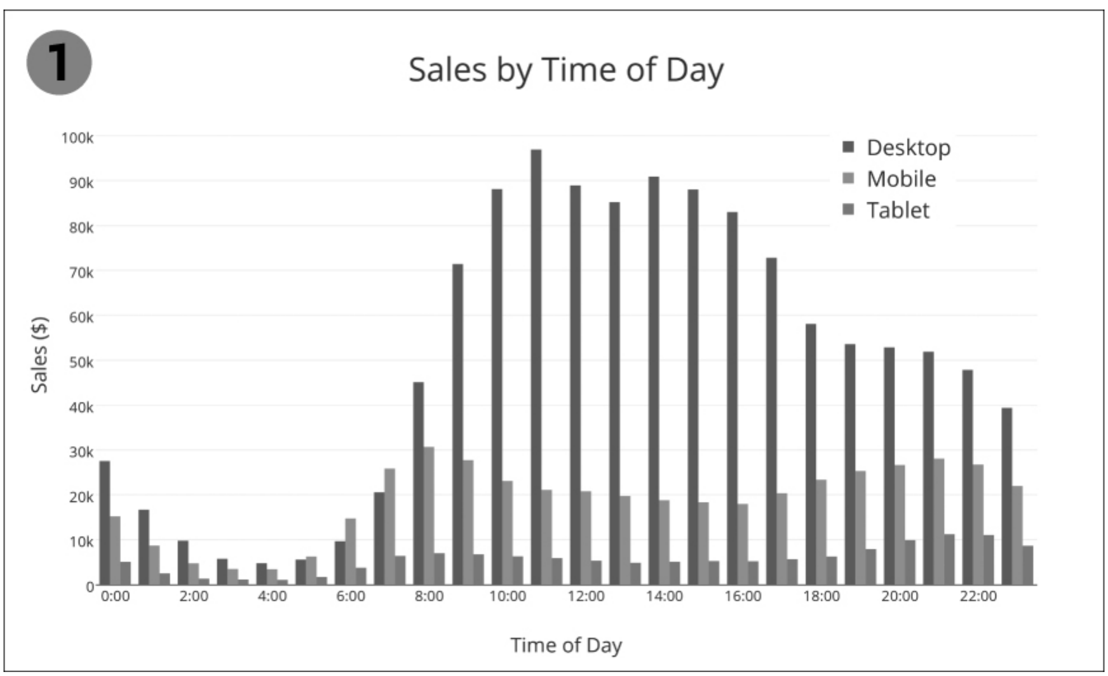

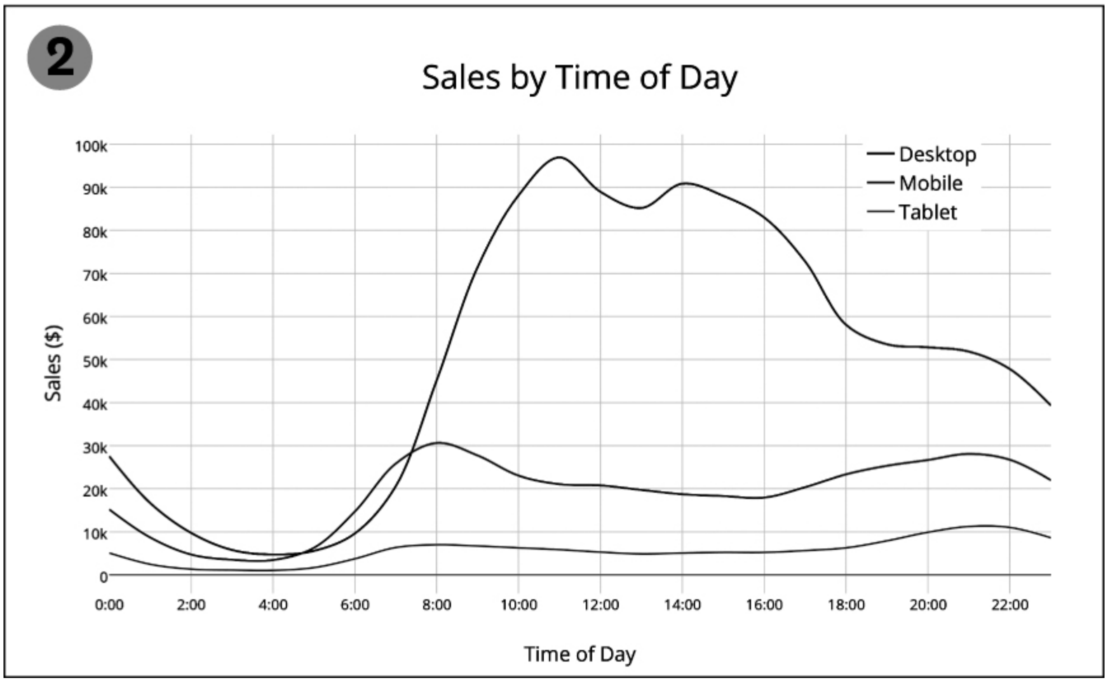

2
Sales by Time of Day
Desktop
Mobile
Tablet
100k
90k
80k
70k
60k
50k
40k
30k
20k
10k
0
Sales ($)
0:00
2:00
4:00
6:00
8:00
10:00
12:00
14:00
16:00
18:00
20:00
22:00
Time of Day

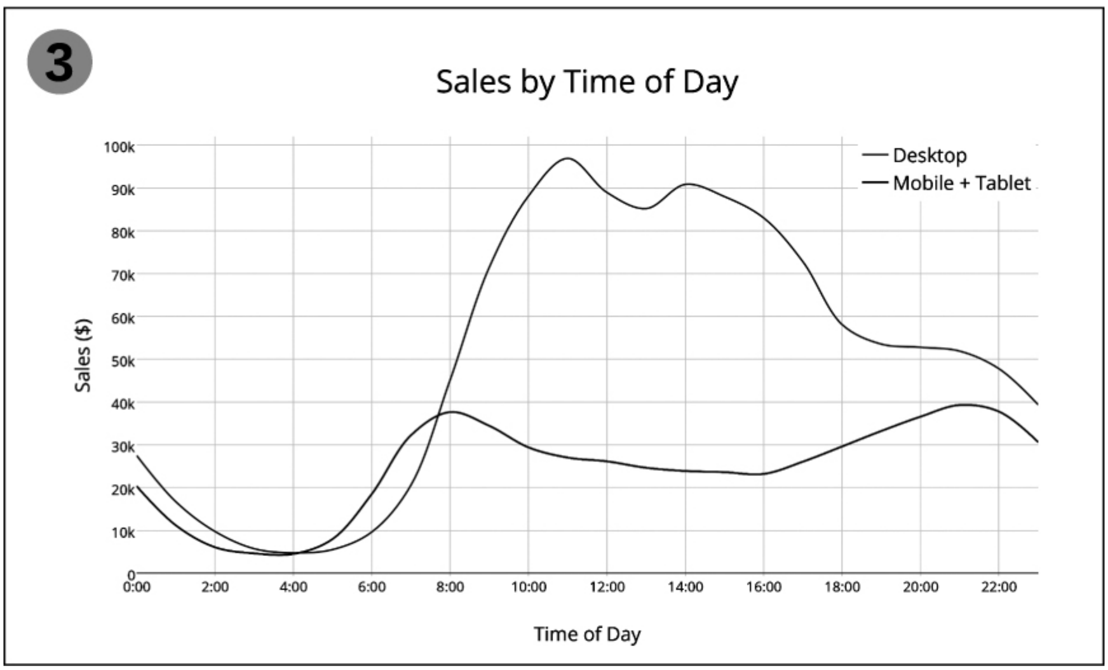

3
Sales by Time of Day
Desktop
Mobile + Tablet
100k
90k
80k
70k
60k
50k
40k
30k
20k
10k
0
Sales ($)
0:00
2:00
4:00
6:00
8:00
10:00
12:00
14:00
16:00
18:00
20:00
22:00
Time of Day

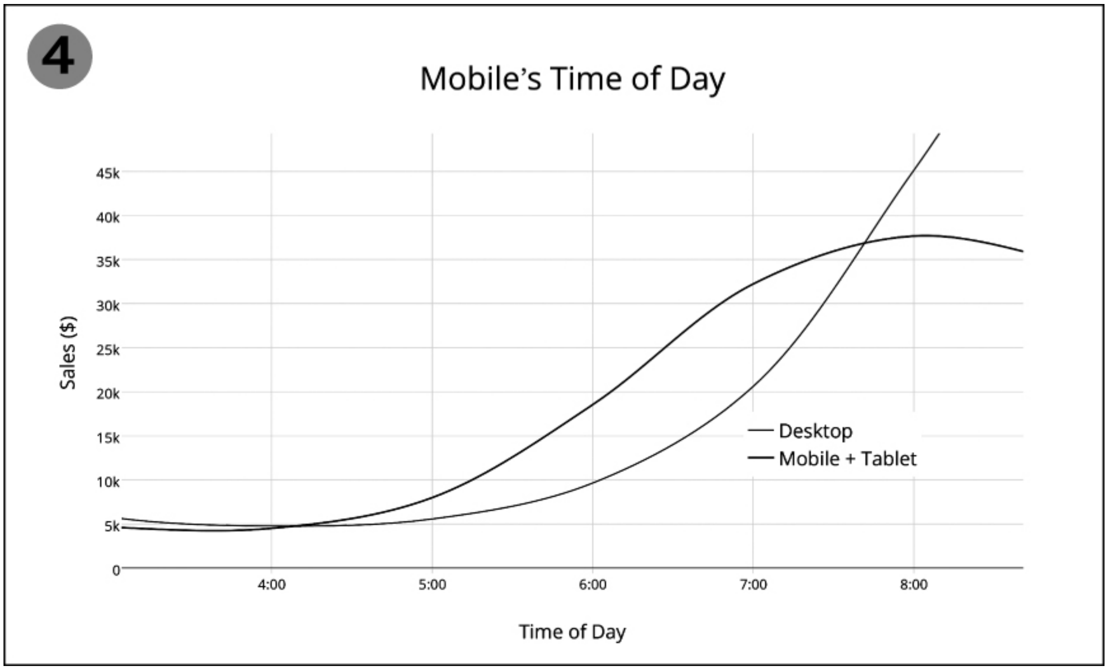

Suppose a marketing manager believes that at certain times of the day more customers shop his site on mobile devices than on desktops, but his marketing programs aren't designed to take advantage of that. He loads some data into an online tool (called Datawrapper) to see if he's right (1 on previous page).

He can't yet confirm or refute his hypothesis. He can't tell much of anything, but he's prototyping and using a tool that makes it easy to try different views into the data. He works fast; design is not a concern. He tries a line chart instead of a bar chart (2).

Now he's seeing something, but working with three variables still doesn't quite get at the mobile-versus-desktop view he wants, so he tries again with two variables (3). Each time he iterates, he evaluates whether he can confirm his original

hypothesis: At certain times of day more customers are shopping on mobile devices than on desktops.

On the fourth try he zooms in and confirms his hypothesis (4).

New software tools mean this type of visualization is easier than ever before: They're making data analysts of us all.

Visual exploration. Open-ended data-driven visualizations tend to be the province of data scientists and business intelligence analysts, although new tools have begun to engage general managers in visual exploration. It's exciting to try, because it often produces insights that can't be gleaned any other way.

Because we don't know what we're looking for, these visuals tend to plot data more inclusively. In extreme cases, this kind of project may combine multiple data sets or load dynamic, real-time data into a system that updates automatically. Statistical modeling benefits from visual exploration.

Exploration also lends itself to interactivity: Managers can adjust parameters, inject new data sources, and continually revisualize. Complex data sometimes also suits specialized and unusual visualization, such as *force-directed diagrams* that show how networks cluster, or topographical plots.

Function trumps form here: Analytical, programming, data management, and business intelligence skills are more crucial than the ability to create presentable charts. Not surprisingly, this half of the quadrant is where managers are most likely to

call in experts to help set up systems to wrangle data and create visualizations that fit their analytic goals.

Anmol Garg, a data scientist at Tesla Motors, has used visual exploration to tap into the vast amount of sensor data the company's cars produce. Garg created an interactive chart that shows the pressure in a car's tires over time. In true exploratory form, he and his team first created the visualizations and then found a variety of uses for them: to see whether tires are properly inflated when a car leaves the factory, how often customers reinflate them, and how long customers take to respond to a low-pressure alert; to find leak rates; and to do some predictive modeling on when tires are likely to go flat. The pressure of all four tires is visualized on a scatter plot, which, however inscrutable to a general audience, is clear to its intended audience.

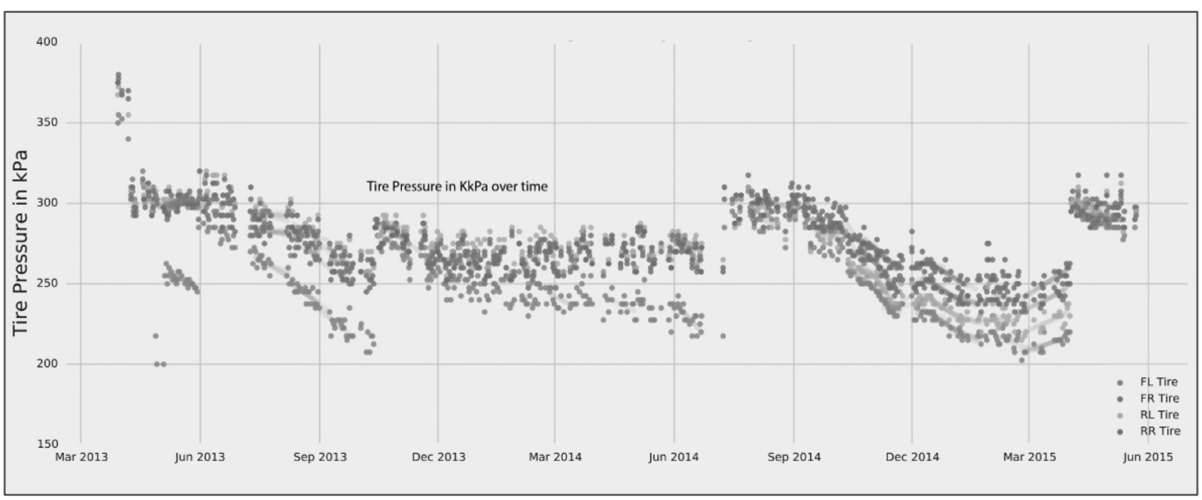

Garg was exploring data to find insights that could be gleaned only through visuals. "We're dealing with terabytes of

data all the time," he says. "You can't find anything looking at spreadsheets and querying databases. It has to be visual." For presentations to the executive team, Garg translates these exploration sessions into the kinds of simpler charts discussed below. "Management loves seeing visualizations," he says.

Everyday dataviz

Info type	Simple, low volume
Typical setting	Formal, presentations
Primary skills	Design, storytelling
Goals	Affirming, setting context

Whereas data scientists do most of the work on visual exploration, managers do most of the work on everyday visualizations. This quadrant comprises the basic charts and graphs you normally paste from a spreadsheet into a presentation. They are usually simple—line charts, bar charts, pies, and scatter plots.

"Simple" is the key. Ideally, the visualization will communicate a single message, charting only a few variables. And the goal is straightforward: affirming and setting context. Simplicity is primarily a design challenge, so design skills are important. Clarity and consistency make these charts most effective in the setting where they're typically used: a formal presentation. In a presentation, time is constrained. A poorly designed chart will waste that time by provoking questions

that require the presenter to interpret information that's meant to be obvious. If an everyday dataviz can't speak for itself, it has failed—just like a joke whose punch line has to be explained.

That's not to say that declarative charts shouldn't generate discussion. But the discussion should be about the idea in the chart, not the chart itself.

Suppose an HR VP will be presenting to the rest of the executive committee about the company's health care costs. She wants to convey that the growth of these costs has slowed significantly, creating an opportunity to invest in additional health care services.

The VP has read an online report about this trend that includes a link to some government data. She downloads the data and clicks on the line chart option in Excel. She has her viz in a few seconds. But because this is for a presentation, she asks a designer colleague to add detail from the data set to give a more comprehensive view.

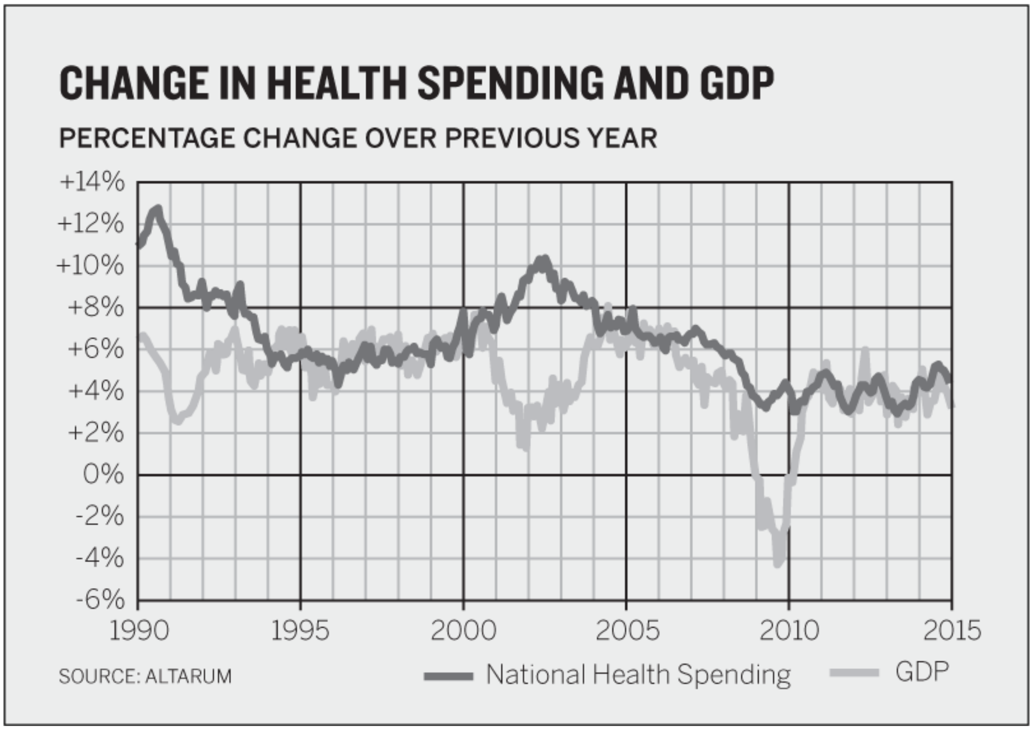

This is a well-designed, accurate chart, but it's probably not the right one. The executive committee doesn't need two decades' worth of historical context to discuss the company's strategy for employee benefits investments. The point the VP wants to make is that cost increases have slowed over the past few years. Is that clearly communicated here?

In general, when it takes more than a few seconds to digest the data in a chart, the chart will work better on paper or on a personal-device screen, for someone who's not expected to listen to a presentation while trying to take in so much information. For example, health care policy makers might benefit from seeing this chart in advance of a hearing at which they'll discuss these long-term trends.

Our VP needs something cleaner for her context. She could make her point as simply as this:

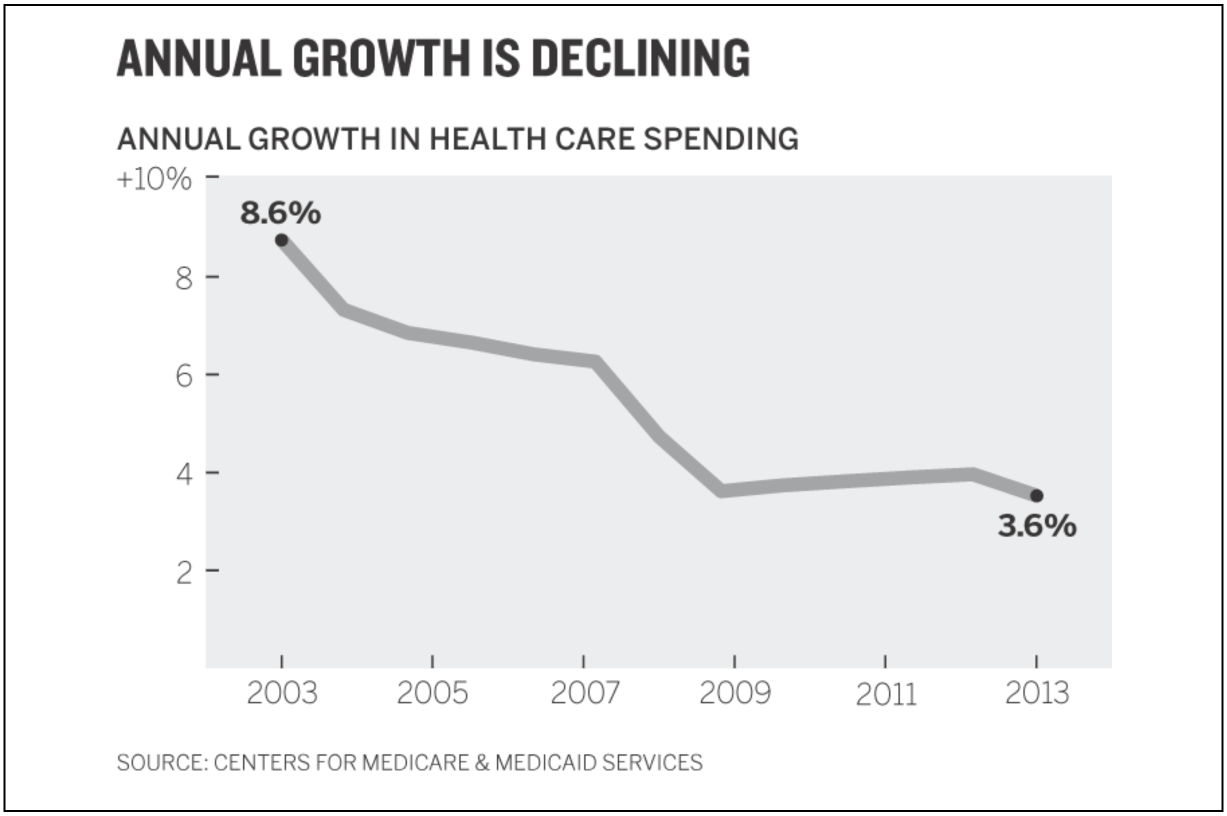

Simplicity like this takes some discipline—and courage—to achieve. The impulse is to include everything you know. Busy charts communicate the idea that you've been just that—busy. "Look at all the data I have and the work I've done," they seem to say. But that's not the VP's goal. She wants to persuade her colleagues to invest in new programs. With this chart, she won't have to utter a word for the executive team to understand the trend. She has clearly established a foundation for her recommendations.

In some ways, “data visualization” is a terrible term. It seems to reduce the construction of good charts to a mechanical procedure. It evokes the tools and methodology required to create rather than the creation itself. It’s like calling *Moby-Dick* a “word sequentialization” or *The Starry Night* a “pigment distribution.”

It also reflects an ongoing obsession in the dataviz world with process over outcomes. Visualization is merely a process. What we actually do when we make a good chart is get at some truth and move people to feel it—to see what couldn’t be seen before. To change minds. To cause action.

Some basic common grammar will improve our ability to communicate visually. But good outcomes require a broader understanding and a strategic approach—which the typology described here is meant to help you develop.

Originally published June 2016. Reprint R1606H

Right Tech, Wrong Time

by Ron Adner and Rahul Kapoor

FOR THE PAST 30 YEARS, "creative destruction" has been a source of fascination at top-tier business schools and in magazines like this one. The almost obsessive interest in this topic is unsurprising, given the ever-changing, never-ending list of transformative threats—which today include the internet of things, 3-D printing, cloud computing, personalized medicine, alternative energy, and virtual reality.

Our understanding of the shifts that disrupt businesses, industries, and sectors has profoundly improved over the past 20 years: We know far more about how to identify those shifts and what dangers they pose to incumbent firms. But the *timing* of technological change remains a mystery. Even as some technologies and enterprises seem to take off overnight (ride sharing and Uber; social networking and Twitter), others take decades to unfold (high-definition TV, cloud computing). For firms and their managers, this creates a problem: Although we have become quite savvy about determining *whether* a new

innovation poses a threat, we have very poor tools for knowing *when* such a transition will happen.

The number-one fear is being ready too late and missing the revolution (consider Blockbuster, which failed because it ignored the shift from video rentals to streaming). But the number-two fear should probably be getting ready too soon and exhausting resources before the revolution begins (think of any dot-com firm that died in the 2001 technology crash, only to see its ideas reborn later as a profitable Web 2.0 venture). This fear of acting prematurely applies both to established incumbents being threatened by disruptive change and to innovating start-ups carrying the flag of disruption.

To understand why some new technologies quickly supplant their predecessors while others catch on only gradually, we need to think about two things differently. First, we must look not just at the technology itself but also at the broader *ecosystem* that supports it. Second, we need to understand that competition may take place *between the new and the old ecosystems,* rather than between the technologies themselves. This perspective can help managers better predict the timing of transitions, craft more-coherent strategies for prioritizing threats and opportunities, and ultimately make wiser decisions about when and where to allocate organizational resources.

You're Only as Good as Your Ecosystem

Both established and disruptive initiatives depend on an array of complementary elements—technologies, services, standards, regulations—to deliver on their value propositions. The strength and maturity of the elements that make up the ecosystem play a key role in the success of new technologies—and the continued relevance of old ones.

The new technology's ecosystem

In assessing an emerging technology's potential, the paramount concern is whether it can satisfy customer needs and deliver value in a better way. To answer that question, investors and executives tend to drill down to specifics: How much additional development will be required before the technology is ready for commercial prime time? What will its production economics look like? Will it be price-competitive?

If the answers suggest that the new technology can really deliver on its promise, the natural expectation is that it will take over the market. Crucially, however, this expectation will hold only if the new technology's dependence on other innovations is low. For example, a new lightbulb technology that can plug into an existing socket can deliver its promised performance right out of the box. In such cases, where the value proposition does not hinge on external factors, great product execution translates into great results.

Idea in Brief

The Problem

Over the past 20 years we've gotten very good at predicting whether a major new technology will supplant an older one—but we are still terrible at predicting when that substitution will take place.

The Insight

If the new technology doesn't need a new ecosystem to support it—if it is essentially plug-and-play—then adoption can be swift. But if other complements are needed, then the pace of substitution will slow until those challenges are resolved. Change takes even longer when the old technology gets a boost from improvements in its own ecosystem.

The Implications

Start-ups need to consider not just when their innovation will be viable, but also what external bottlenecks will arise. Incumbents, meanwhile, should use the transition period to up their own game—and to figure out a strategy for long-term survival.

However, many technologies do not fall into this plug-and-play mold. Rather, their ability to create value depends on the development and commercial deployment of other critical parts of the ecosystem. Consider HDTV, which could not gain traction until high-definition cameras, new broadcast standards, and updated production and postproduction processes also became commercially available. Until the entire

ecosystem was ready, the technology revolution promised by HDTV was bound to be delayed, no matter how great its potential for a better viewing experience. For the pioneers who developed HDTV technology in the 1980s, being right about the vision brought little comfort during the 30 years it took for the rest of the ecosystem to emerge.

An improved lightbulb and an HDTV both depend on ecosystems of complementary elements. The difference is that the lightbulb plugs into an existing ecosystem (established power generation and distribution networks; wired homes), whereas the television requires the successful development of co-innovations. Improvements in the lightbulb will thus create immediate value for customers, but the TV's ability to create value is limited by the availability and progress of other elements in its ecosystem.

The old technology's ecosystem

Successful, established technologies—by definition—have overcome their emergence challenges and are embedded within successful, established ecosystems. Whereas new technologies can be held back by their ecosystems, incumbent technologies can be accelerated by improvements in theirs, even in the absence of progress in the core technology itself. For example, although the basic technology behind bar codes has not changed in decades, their utility improves every year

as the IT infrastructure supporting them allows ever-more information to be extracted. Hence in the 1980s, bar codes allowed prices to be automatically scanned into cash registers; in the 1990s, aggregating the bar code data from daily or weekly transactions provided insight into general inventory; in the modern era, bar code data is used for real-time inventory management and supply chain restocking. Similarly, improvements in DSL (digital subscriber line) technology have extended the life of copper telephone lines, which can now offer download speeds of 15 megabytes per second, making copper-wire services competitive with newer cable and fiber networks.

The War Between Ecosystems

When a new technology isn't a simple plug-and-play substitution—when it requires significant developments in the ecosystem in order to be useful—then a race between the new- and the old-technology ecosystems begins.

What determines who wins? For the *new* technology, the key factor is how quickly its ecosystem becomes sufficiently developed for users to realize the technology's potential. In the case of cloud-based applications and storage, for example, success depended not just on figuring out how to manage data in server farms, but also on ensuring the satisfactory performance of critical complements such as broadband and

online security. For the *old* technology, what's important is how its competitiveness can be increased by improvement in the established ecosystem. In the case of desktop storage systems (the technology that cloud-based applications would replace), extension opportunities have historically included faster interfaces and more-robust components. As these opportunities become exhausted, we can expect substitution to accelerate.

About the Research

WE DEVELOPED AND EXPLORED the ideas described in this article during a five-year research project on the pace of substitution in the semiconductor-manufacturing ecosystem.

The semiconductor industry's remarkably robust progress over the past 60 years was made possible by innovations in the lithography technology that semiconductor manufacturers use. We studied the successive generations of lithography equipment and noticed a pattern: In some cases, the new technology dominated the market in a matter of two to five years, whereas in other cases it faced prolonged, unexpected delays in achieving market dominance—and sometimes never did. This was true despite the fact that each generation offered superior performance, even on a price-adjusted performance basis.

To test our hypotheses about how ecosystem emergence challenges and extension opportunities affect the pace of

substitution, we first collected and analyzed detailed data on every product and firm involved in every generation of the technology. We supplemented that information with extensive interviews with executives from firms throughout the ecosystem.

Our statistical analysis showed that 48% of the variation in the pace of substitution was attributable to traditional factors: price-adjusted performance differences, the number of rivals in the market, and the tenure of the old technology. When we added consideration of the ecosystem dynamics discussed in the article, we were able to account for a remarkable 82% of the variance.

For more details on the research, see "Innovation Ecosystems and the Pace of Substitution: Re-examining Technology S-Curves," by Ron Adner and Rahul Kapoor, *Strategic Management Journal* (March 2015).

Thus the pace of substitution is determined by the rate at which the new technology's ecosystem can overcome its emergence challenges relative to the rate at which the old technology's ecosystem can exploit its extension opportunities. To consider the interplay between these forces, we have developed a framework to help managers assess how quickly disruptive change is coming to their industry (see the chart "A framework for analyzing the pace of technology substitution"). There are four possible scenarios: creative destruction, robust resilience, robust coexistence, and the illusion of resilience.

Creative destruction

When the ecosystem emergence challenge for the new technology is low and the ecosystem extension opportunity for the old technology is also low (quadrant 1 in the framework), the new technology can be expected to achieve market dominance in short order (see point A in the exhibit "How fast does new technology replace the old?"). The new technology's ability to create value is not held back by bottlenecks elsewhere in the ecosystem, and the old technology has limited potential to improve in response to the threat. This quadrant aligns with concept of creative destruction—the idea that an innovative upstart can swiftly cause the demise of established competitors. While the old technology can continue serving niches for a long time (see "Bold Retreat," by Ron Adner and Daniel C. Snow, HBR, March 2010), the bulk of the market will abandon it relatively quickly in favor of the new technology. As an example, consider the rapid replacement of dot matrix printers by inkjet printers.

Robust resilience

When the balance is reversed—when the new technology's ecosystem confronts serious emergence challenges and the old technology's ecosystem has strong opportunities to improve (quadrant 4)—the pace of substitution will be very slow. The old technology can be expected to maintain a prosperous

leadership position for an extended period. This quadrant is most consistent with technologies that seem revolutionary when they're first touted but appear overhyped in retrospect.

A framework for analyzing the pace of technology substitution

The pace of substitution is determined by how quickly the new technology's ecosystem challenges are resolved and whether the old technology can exploit ecosystem opportunities for extension.

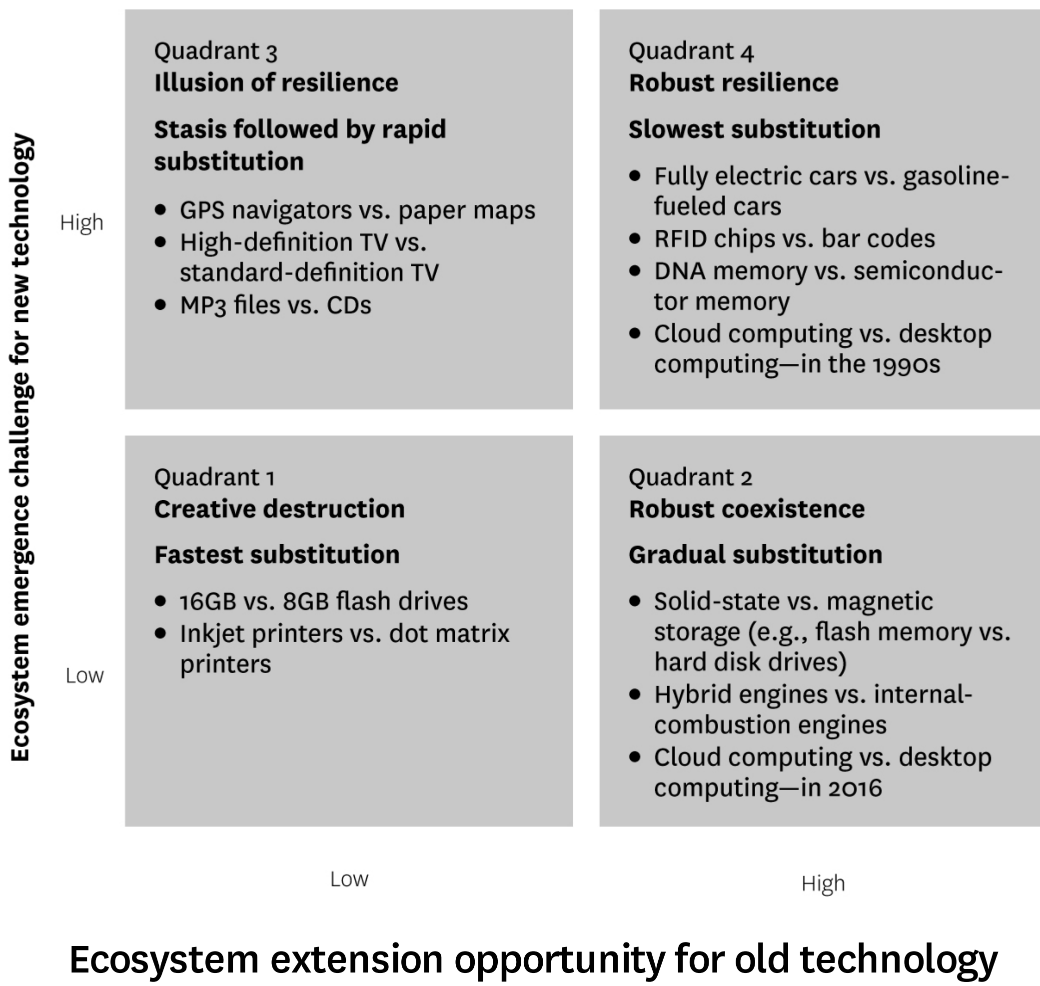

Bar codes and radio frequency identification (RFID) chips provide a good example. RFID chips hold the promise of storing far richer data than bar codes ever could, but their adoption has lagged because of the slow deployment of suitable IT infrastructure and nonuniform industry standards. Meanwhile, IT improvements have extended the usability of bar code data, as we've already discussed, relegating RFID to niche applications and keeping the RFID revolution at bay for the past two decades. It may well be that RFID does eventually overcome its challenges and that ecosystem extension opportunities dry up for bar codes. If this happens, the dynamics will shift from quadrant 4 to another quadrant, and the pace of substitution will quicken. But that will be small consolation to the firms and investors that committed to RFID decades ago. The opportunity cost of waiting for the rest of the system to catch up can mean that being in the right place 10 years too soon is more costly than missing the revolution completely.

How fast does new technology replace the old?

Traditionally the substitution of a new technology for an old one is shown with two S curves (the solid lines). A more holistic view adds two more dynamics. First, if the new technology depends on the emergence of a new ecosystem, it becomes dominant more slowly (dotted line). Second, the old technology's competitiveness is

extended if it can benefit from performance improvements in its surrounding ecosystem (dashed line).

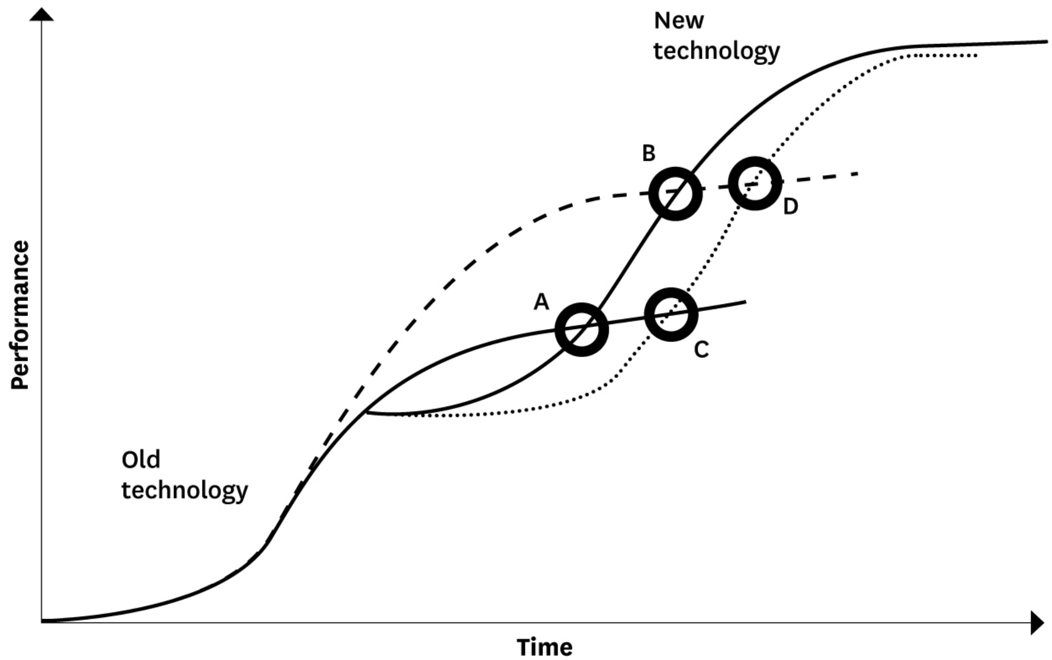

Creative destruction

Point A
The classic—and fastest—substitution takes place when the new technology's ecosystem is ready to go and the old technology's ecosystem can't be significantly improved.

Robust coexistence

Point B
If the new technology is compatible with the existing ecosystem and the old technology's ecosystem can be significantly improved, substitution takes place later (relative to creative destruction) and at a higher performance level.

Illusion of resilience

Point C
If the new technology's ecosystem needs considerable development and the old technology's ecosystem has little room for improvement, the changeover occurs after time has passed without performance gains.

Robust resilience

Point D
If the new technology's ecosystem needs considerable development and there are abundant opportunities to improve the old technology's ecosystem, the substitution occurs after the longest period of time and at the highest performance level.

Note: The exact positions of B and C will depend on the specifics of the case, but they will reflect an intermediate pace of substitution (relative to points A and D) and intermediate performance at substitution.

When substitution is slow, there are also implications for the new technology's required performance levels (see point D in the exhibit). Every time IT improvements make bar codes more useful, for example, the quality threshold for the RFID technology is raised. Thus performance expectations for the innovation keep ratcheting upward, even as its wide adoption is held back by the underdeveloped state of its ecosystem.

Robust coexistence

When the ecosystem emergence challenge for the new technology is low and the ecosystem extension opportunity for the old technology is high (quadrant 2), competition will be robust. The new technology will make inroads into the market, but improvements in the old-technology ecosystem will allow the incumbent to defend its market share. There will be a prolonged period of coexistence. Although extension opportunities are unlikely to reverse the rise of the new technology, they will materially delay its dominance.

An instructive example is the competition between hybrid (gas-electric) automobile engines and traditional internal-combustion engines. Unlike fully electric engines, which need a supporting network of charging stations, hybrids were not held back by ecosystem emergence challenges. At the same time, however, traditional gas engines have become more fuel-efficient, and the ecosystem for the traditional technology has

improved, too, as gas engines have become better integrated with other elements in the vehicle, such as heating and cooling systems.

A period of robust coexistence can be quite attractive from a consumer perspective. Performance of both ecosystems is improving—and the better the old technology's ecosystem becomes, the higher the performance bar is for the new technology's ecosystem (point B in the exhibit).

The illusion of resilience

When the ecosystem emergence challenge is high for the new technology and the ecosystem extension opportunity is low for the old technology (quadrant 3), not much will change until the emergence challenge is resolved—but then substitution will be rapid (point C in the exhibit). Examples here are HDTVs versus traditional TVs, and e-books versus printed books. Both of those revolutions were delayed not by advances in the old technology's ecosystem but by ecosystem-emergence challenges in the new technology.

In scenarios in this quadrant, an industry analysis will most likely show that the old technology maintains high market share, but growth has stalled. Because rapid market-share inversion is to be expected once the new technology fulfills its value creation potential, the dominance of the old technology

is fragile. It is maintained not by continued progress in the old technology but by setbacks for the new competitor.

Implications for Action

Once you understand that in the race to dominance, ecosystems are just as important as technologies, you will be better at thinking through how quickly change is going to occur—and deciding what level of performance you need to aim for in the meantime. We will consider how to tackle these questions shortly, but first let's review a few general truths that emerge from this perspective.

- If your company is introducing a potentially transformative innovation, the full value will not be realized until all bottlenecks in the ecosystem are resolved. It may pay to focus a little less on perfecting the technology itself and a little more on resolving the most pressing problems in the ecosystem.
- If you are a threatened incumbent, it pays to analyze not just the emerging technology itself but also the ecosystem that supports it. The greater the ecosystem-emergence challenge for the new technology, the more time you have to strengthen your own performance.
- Strengthening incumbent performance may mean improving the old technology—but it can just as easily

mean improving aspects of the ecosystem that supports it.

- Every time the old technology's performance gets better, the performance threshold for the new technology goes up.

With that overview in mind, let's look at how to use this framework to analyze your own technology strategy. We recommend having executive conversations focused on two questions: Which quadrant is our industry in? and What are the implications for our resource allocation and other strategic choices?

Which quadrant are we in?

Without the benefit of hindsight, your response to this question is clearly a matter of judgment. Some people would look at electric vehicles in 2016 and say they are still stuck in quadrant 4 (where we have placed them in our framework), pointing out that the charging infrastructure and battery performance are insufficient for mainstream adoption. Other people would position EVs on the cusp of quadrant 2, claiming that acceptance is growing and that better batteries make it possible to drive longer distances before recharging. Still others would place EVs solidly in quadrant 2, arguing that Tesla's success in selling its vehicles and populating its waiting lists is a sure sign that commercial potential is no longer constrained.

The sidebar "How Big a Threat Is the New Technology?" suggests issues to think through as you debate which quadrant

you're in. Some questions pertain to the new technology and some to the old—but you will want to consider them all, regardless of whether you are an incumbent or a start-up. Don't expect all individual team members to agree on the answers to these questions. It is precisely by going through the process of articulating different views that teams can make the most of their collective insights.

How Big a Threat Is the New Technology?

PREDICTING THE PACE OF SUBSTITUTION requires analyzing the competition between the new- and the old-technology ecosystems. Six questions can help innovators and incumbents assess their positions and strategies.

New-Technology Questions

These questions (drawn from *The Wide Lens*, by coauthor Ron Adner) address the emergence challenges that confront the new technology. The answers should help innovators decide how to adjust their strategies.

1. What is the *execution risk*—the level of difficulty in delivering the focal innovation to the market on time and to spec?
2. What is the *co-innovation risk*—the extent to which the success of the new technology depends on the successful commercialization of other innovations?

3. What is the *adoption-chain risk*—the extent to which other partners need to adopt and adapt to the new technology before end consumers can fully assess its value proposition?

The greater the extent to which the new technology is facing any of these risks, the greater the challenge to be overcome, and the longer the expected delay in adoption of the technology.

Old-Technology Questions

These questions address the prospects for improving the competitiveness of the incumbent technology. The answers should help incumbents identify opportunities they might exploit.

1. Can the competitiveness of the old technology be extended by further improvements to the technology itself?
2. Can it be extended by improvements to complementary elements in its ecosystem?
3. Can it be extended by borrowing from innovations in the new technology and its ecosystem?

The more positive the reply to each of these questions, the greater the extension opportunity for the old technology.

What are the implications for resource allocation and other strategic choices?

Each quadrant in the framework carries different implications for resource allocation decisions. And since markets are not transformed all at once, the quadrant also suggests possible ways to position yourself during the transition.

In quadrant 1 (creative destruction), with the old technology stagnant and the new technology unhampered, innovators should aggressively invest in the new technology. Incumbents should follow the familiar prescriptions for embracing change to withstand the winds of creative destruction. Part of that is looking for niche positions where they can survive in the long term with the old technology. For example, pagers were largely replaced by cell phones, but they are still used by emergency-service providers.

In quadrant 2 (robust coexistence), incumbent firms can continue to invest in the old technology and aggressively invest in improvements to the ecosystem, knowing that the new and the old technologies will coexist for an extended period. As in quadrant 1, they should also seek niche positions for the old technology for the long term, but there is less urgency to do so. New-technology innovators should move full speed ahead on perfecting the new technology along with its complements. That includes testing and refining the offering with early adopters and segments that are potentially receptive.

In quadrant 3 (the illusion of resilience), new-technology champions should direct resources toward resolving their

ecosystem challenges and developing complementary elements, and resist overprioritizing further development of the technology itself. When the bottleneck to adoption is the ecosystem, not the technology, pushing technology progress is pushing the wrong lever. Incumbents, for their part, must guard against the false assumption that they are maintaining their market position because of the merits of their own technology. As publishers of road atlases will attest, this is probably a time to harvest and make only incremental improvements, with an eye toward sunset; it is not the time to redouble innovation efforts in the old technology.

Finally, in quadrant 4 (robust resilience), incumbent firms should invest aggressively in upgrading their offerings and actively raising the bar that challengers need to cross. Obviously, new-technology innovators should be clear-eyed about working to resolve the ecosystem constraints they face. But at the same time they must recognize that the performance threshold for their core technology is rising. That necessitates both a significant level of resource investment and considerable patience regarding investment returns. Innovators are not likely to transform the sector in the foreseeable future, and therefore they will want to think through the economics of serving those customers they can succeed with.

One final note about the dynamics of change. Every innovator wants to end up in quadrant 1 so that it can play the

classic creative-destruction game. But there are different paths for getting there. A hypothesis that predicts a transition path from Q4 to Q3 to Q1 is a bet on the exhaustion of the old technology. For an innovator, that would mean focusing on aligning the new-technology ecosystem without great concern for extending a performance advantage. In contrast, a predicted path of Q4 to Q2 to Q1 would mean competing against an improving incumbent-technology ecosystem. Here the innovator needs to continually elevate its performance while it simultaneously perfects the ecosystem.

Few modern firms are untouched by the urgency of innovation. But when it comes to strategizing for a revolution, the question of "whether" often drowns out the question of "when." Unfortunately, getting the first right but not the second can be devastating. "Right tech, wrong time" syndrome is a nightmare for any innovating firm. Closer analysis of the enabling contexts of rival technologies—Is the new ecosystem ready to roll? Does the old ecosystem still hold potential for improvement?—sheds more light on the question of timing. And better timing, in turn, will improve the efficiency and effectiveness of the innovation efforts that are so critical for survival and success.

Further Reading

FOR MORE INSIGHTS INTO THE relationship between technologies and their ecosystems, see the following:

- **Match Your Innovation Strategy to Your Innovation Ecosystem,"** Ron Adner, HBR, April 2006
- **A Sad Lesson in Collaborative Innovation,"** Ron Adner, HBR.org, May 9, 2012
- ***The Wide Lens: What Successful Innovators See That Others Miss,*** Ron Adner, Portfolio/Penguin 2013
- **Beware of Old Technologies' Last Gasps,"** Daniel Snow, HBR, January 2008
- **Value Creation in Innovation Ecosystems: How the Structure of Technological Interdependence Affects Firm Performance in New Technology Generations,"** Ron Adner and Rahul Kapoor, *Strategic Management Journal* March 2010

Originally published in November 2016. Reprint R1611C

How to Pay for Health Care

by Michael E. Porter and Robert S. Kaplan

THE UNITED STATES STANDS at a crossroads as it struggles with how to pay for health care. The fee-for-service system, the dominant payment model in the U.S. and many other countries, is now widely recognized as perhaps the single biggest obstacle to improving health care delivery.

Fee for service rewards the quantity but not the quality or efficiency of medical care. The most common alternative payment system today—fixed annual budgets for providers—is not much better, since the budgets are disconnected from the actual patient needs that arise during the year. Fixed budgets inevitably lead to long waits for nonemergency care and create pressure to increase budgets each year.

We need a better way to pay for health care, one that rewards providers for delivering superior value to patients: that is, for achieving better health outcomes at lower cost. The move toward "value-based reimbursement" is accelerating, which is an encouraging trend. And the Centers for Medicare &

Medicaid Services (CMS), to its credit, is leading the charge in the United States.

That doesn't mean, however, that health care is converging on a solution. The broad phrase "value-based reimbursement" encompasses two radically different payment approaches: capitation and bundled payments. In capitation, the health care organization receives a fixed payment per year per covered life and must meet all the needs of a broad patient population. In a bundled payment system, by contrast, providers are paid for the care of a patient's medical condition across the entire care cycle—that is, all the services, procedures, tests, drugs, and devices used to treat a patient with, say, heart failure, an arthritic hip that needs replacement, or diabetes. If this sounds familiar, it's because it is the way we usually pay for other products and services we purchase.

A battle is raging, largely unbeknownst to the general public, between advocates of these two approaches. The stakes are high, and the outcome will define the shape of the health care system for many years to come, for better or for worse. While we recognize that capitation can achieve modest savings in the short run, we believe that it is not the right solution. It threatens patient choice and competition and will fail to fundamentally change the trajectory of a broken system. A bundled payment system, however, would truly transform the way we deliver care and finally put health care on the right path.

The Small Step: Capitation

Capitation, or population-based payment, is not a new idea. It was introduced in the United States with some fanfare in the 1990s but quickly ran into widespread criticism and was scaled back significantly. Today, a number of transitional approaches, including accountable care organizations (ACOs), shared savings plans, and alternative quality contracts, have been introduced as steps toward capitation. In the ACO model, the care organization earns bonuses or penalties on the basis of how the total fee-for-service charges for all the population's treatments during the year compare with historical charges. In full capitation, the care organization absorbs the difference between the sum of capitation payments and its actual cost.

Under capitation, unlike in the FFS model, the payer (insurer) no longer reimburses various providers for each service delivered. Rather, it makes a single payment for each subscriber (usually per patient per month) to a single delivery organization. The approach rewards providers for lowering the overall cost of treating the population, which is a step forward. However, under this system cost reduction gravitates toward population-level approaches targeting generic high-cost areas, such as limiting the use of expensive tests and drugs, reducing readmissions, shortening lengths of stay, and discharging patients to their homes rather than to higher-cost rehabilitation facilities. As a response to the failed experience

with capitation in the 1990s, current capitation approaches include some provider accountability for quality. However, “quality” is measured by broad population-level metrics, such as patient satisfaction, process compliance, and overall outcomes such as complication and readmission rates.

Idea in Brief

The Challenge

The United States stands at a crossroads as it struggles with how to pay for health care. Fee for service, the dominant model today, is widely recognized as the single biggest obstacle to improving health care delivery. The choice is between two fundamentally different approaches: capitation and bundled payments. The stakes are high, and the outcome will define the shape of the health care system for many years to come, for better or for worse.

The Danger

Although capitation may deliver modest savings in the short run, it is not the solution. It entrenches large existing systems, eliminates patient choice, promotes consolidation, limits competition, and perpetuates the lack of accountability for outcomes. Like fee for service, capitation will fail to drive true innovation in health care delivery.

The Opportunity

Bundled payments trigger competition among providers to create value where it matters—at the individual patient level

—and will finally put health care on the right path. Robust proof-of-concept initiatives in the U.S. and abroad demonstrate that the challenges of transitioning to bundled payments are already being overcome.

This all seems good at first blush. The trouble is that, like the failed FFS payment system, capitation creates competition at the wrong level and on the wrong things, rather than on what really matters to patients and to the health care system overall.

Providers are not accountable for patient-level value

Capitation and its variants reward improvement at the population level, but patients don't care about population outcomes such as overall infection rates; they care about the treatments they receive to address their particular needs. Outcomes that matter to breast cancer patients are different from those that are important to patients with heart failure. Even for primary and preventive care, which the concept of population health rightly emphasizes, appropriate care depends heavily on each patient's circumstances—health status, comorbidities, disability, and so on. And managing the overall health of a diverse population with high turnover (as ACOs do) is extremely difficult.

Thus, capitated payments are not aligned with better or efficient care for each patient's particular condition. Instead, capitation puts the focus on limiting the overall amount of care delivered without tying the outcomes back to individual

patients or providers. The wrong incentives are created, just as is the case for fee for service, which reimburses for the volume of services but not the value.

Providers bear the wrong risks

Because capitation pays providers a fee per person covered, it shifts the risk for the cost of the population's actual mix of medical needs—over which they have only limited control—to providers. Some large private insurers favor capitation for just this reason. But bearing the actuarial risk of a population's medical needs is what insurers should do, since they cover a far larger and more diverse patient population over which to spread this risk. Providers should bear only the risks related to the actual care they deliver, which they can directly affect.

A more fundamental problem is that capitation payments are extremely difficult to adjust to reflect each patient's overall health risk, not to mention to correctly adjust for this risk across a large, diverse population. Risks are much better understood and managed for a particular medical condition—for example, the probable effects of age or comorbidities on the costs and outcomes for joint replacement—as is the case in bundled payments.

Because population-level risk factors are so complex, health systems under capitation have an incentive to claim as many comorbidities as possible to bolster their revenue and profitability. A whole segment of health care IT providers has

emerged to help providers “upcode” patients into higher-risk categories. Such gaming of risk adjustment first became a problem during the era of managed-care capitation in the 1990s, and it remains one today.

Patient choice is limited, and competition is threatened

Capitation creates strong incentives for a health system to deliver all the care within its system, because contracting for outside services reduces net revenue and results in underutilization of existing internal capacity. There is even a term for this in health care—“avoiding leakage”—and many systems explicitly monitor and control it. Capitated health systems encourage or require patients (and their referring doctors) to use in-house providers (the ultimate narrow network). Patients are often penalized with extra fees when they don’t use services within the system, even if outside providers have greater experience and get better results for treating the patient’s particular condition. Capitation creates, in essence, a monopoly provider for all the patients in the population. Consumers cannot choose the best provider for their particular needs.

Since providers now bear actuarial risk, they also have a strong incentive to amass the largest possible population. This will accelerate the recent trend of providers’ buying up other hospitals and physician practices and merging systems, which reduces competition. To offset health systems’ rising

bargaining power, insurers will feel pressure to merge. The two dynamics will reinforce each other as provider consolidation begets even more insurer consolidation.

The end result will be the emergence of a few dominant systems—or even only one—in each region. This would be bad for patients. No one organization can have all the skills and technologies needed to be the best in treating everything. We need multiple providers in each region to ensure enough choice and drive innovation in care delivery.

The bottom line is that capitation is the wrong way to pay for health care. It is a top-down approach that achieves some cost savings by targeting low-hanging fruit such as readmission rates, expensive drugs, and better management of post-acute care. But it does not really change health care delivery, nor does it hold providers accountable for efficiency and outcomes where they matter to patients—in the treatment of their particular condition. Capitation's savings also come at the high cost of restricting patient choice and inhibiting provider competition.

Let's consider the alternative.

Paying for Value: Bundled Payments

For virtually all types of products and services, customers pay a single price for the whole package that meets their needs. When purchasing a car, for example, consumers don't buy the

motor from one supplier, the brakes from another, and so on; they buy the complete product from a single entity. It makes just as little sense for patients to buy their diagnostic tests from one provider, surgical services from another, and post-acute care from yet another. Bundled payments may sound complicated, but in setting a single price for all the care required to treat a patient's particular medical condition, they actually draw on the approach long used in virtually every other industry.

Bundled payments have existed in health care for some time in isolated fields such as organ transplantation. They are also common for services that patients pay for directly, such as Lasik eye surgery, plastic surgery, and in vitro fertilization.

To maximize value for the patient, a bundled payment must meet five conditions:

Payment covers the overall care required to treat a condition

The bundled payment should cover the full cost of treating a patient over the entire care cycle for a given condition or over time for chronic conditions or primary care. The scope of care should be defined from the patient's perspective ("Delivering a healthy child"). Care should include all needed services, including managing common comorbidities and related complications. In primary and preventive care, bundled payments should include all the needed care for each defined patient segment (such as healthy adults or low-income elderly).

Payment is contingent on delivering good outcomes

Bundled payments should be tied to achieving the outcomes that matter to patients for each condition and primary care patient segment. Important outcomes include maintaining or returning to normal function, reducing pain, and avoiding and reducing complications or recurrences.

Payment is adjusted for risk

Differences in patients' age and health status affect the complexity, outcomes, and cost of treating a particular condition, as do their social and living circumstances. These risk factors should be reflected in the bundled payment and in expectations for outcomes to reward providers for taking on hard cases.

Payment provides a fair profit for effective and efficient care

A bundled payment should cover the full costs of the necessary care, plus a margin, for providers that use effective and efficient clinical and administrative processes. It should not cover unnecessary services or inefficient care.

Providers are not responsible for unrelated care or catastrophic cases

Providers should be responsible only for care related to the condition—not for care such as emergency treatment after an accident or an unrelated cardiac event. The limits of provider responsibility should be specified in advance and subject to

adjudication if disputes arise. Bundled payments should also include a “stop loss” provision to limit providers’ exposure to unusually high costs from catastrophic or outlier cases. This reduces the need for providers to build such costs into the price for every patient (unlike in capitation).

How Bundled Payments Will Transform Patient Care

Decades of incremental efforts to cut costs in health care and impose practice guidelines on clinicians have failed. Bundled payments directly reward providers for delivering better value for the patient’s condition and will unlock the restructuring of health care delivery in three crucial ways that capitation cannot.

Integrated, multidisciplinary care

Specialty silos have historically led to fragmented, uncoordinated, and inefficient care. With bundled payments, providers with overall responsibility for the full care cycle for a condition will be empowered and motivated to coordinate and integrate all the specialists and facilities involved in care. Clinical teams (the experts) have the freedom to decide how to spend the fixed bundled payment, rather than being required to deliver the services that are reimbursed by legacy FFS payments in order to receive revenue. Teams can choose to add services that are not currently covered by FFS but that provide value for patients.

Bundled payments are triggering a whole new level of care innovation. For example, hospital-based physicians are remaining involved in care after patients are discharged. Hospitalists are added to teams to coordinate all the inpatient specialists involved in the care cycle. Nurses make sure patients fill their prescriptions, take medications correctly, and actually see their primary care physician. (A recent study showed that 50% of readmitted patients did not see their primary care doctor in the first 30 days after discharge.) And navigators accompany patients through all phases of their care and act as first responders in quickly resolving problems. Bundled payments are also spurring innovation in the creation of tailored facilities, such as those of Twin Cities Orthopedics (Minneapolis), which performs joint-replacement care in outpatient surgery centers and nearby recovery centers, rather than in a traditional hospital.

Bundled payments will accelerate the formation of integrated practice units (IPUs), such as MD Anderson's Head and Neck Center and the Joslin Diabetes Center. IPUs combine all the relevant clinicians and support personnel in one team, working in dedicated facilities. Joslin, for example, brings together all the specialists (endocrinologists, nephrologists, internists, neurologists, ophthalmologists, and psychiatrists) and all the support personnel (nurses, educators, dieticians, and exercise physiologists) required to provide high-value diabetes care. IPUs concentrate volume of patients with a given

condition in one place, allowing diagnosis and treatment by a highly experienced team. Numerous studies show that this approach leads to better outcomes and greater efficiency (including less wait time and fewer visits). Bundled payments also encourage the formation of “virtual” IPUs, where even separate practices and organizations actively collaborate across inpatient and outpatient settings to coordinate and integrate care—something that rarely happens today.

How Fee for Service Destroys Value for Patients

FEE-FOR-SERVICE REIMBURSEMENT, the dominant method used to pay for health care in the United States and elsewhere, has held back improvements in the quality of care and led to escalating costs. Overturning the status quo is not easy, but here’s why doing so is essential.

Rewards poor outcomes

Because FFS reimburses providers on the basis of volume of care, providers are rewarded not just for performing unnecessary services but for poor outcomes. Complications, revisions, and recurrences all result in the need for additional services, for which providers get reimbursed again.

Fosters duplication and lack of coordination

FFS makes payments for individual procedures and services, rather than for the treatment of a patient’s condition over the

entire care cycle. In response, providers have organized around functional specialties (such as radiology). Today, multiple independent providers are involved in each patient's treatment, resulting in poorly coordinated care, duplicated services, and no accountability for health outcomes.

Perpetuates inefficiency

Today's FFS payments reflect historical reimbursements with arbitrary inflation adjustments, not true costs. Reimbursement levels vary widely, causing cross-subsidization across specialties and particular services. The misalignment means that inefficient providers can survive, and even thrive, despite high costs and poor outcomes.

Reduces focus

FFS motivates providers to offer full services for all types of conditions to grow overall revenue, even as internal fragmentation causes patients to be handed off from one specialty to another. By attempting to cater to a diverse population of patients, providers fail to develop the specialized capabilities and experience in any one condition necessary for the delivery of excellent care.

Accountability for outcomes

By definition, a bundled payment holds the entire provider team accountable for achieving the outcomes that matter to patients for their condition—unlike capitation, which involves only loose accountability for patient satisfaction or population-level quality targets.

Because bundled payments are adjusted for risk, providers are rewarded for taking on difficult cases. With a fixed single payment, they are penalized if they overtreat patients or perform care in unnecessarily high-cost locations. And because providers are accountable for outcomes covering the entire care cycle, they will move quickly to add new services, more-expensive interventions, or better diagnostic tests if those will improve outcomes or lower the overall cost of care. Specialists operating under a bundled payment, for example, have added primary care physicians to their care teams to better manage the overall care cycle and deal with comorbidities.

Most important, the accountability built into bundled payments will finally bring to health care the systematic measurement of outcomes at the condition level, where it matters most. We know from every other field that measuring and being accountable for results is the most powerful driver of innovation and continuous improvement.

Cost reduction

There have been repeated efforts to control health costs for decades without success, and top-down cost reduction initiatives have sometimes increased costs rather than reduced them. The core problem is that legacy payment models such as FFS have given providers no incentive to cut costs or even to understand what their costs are for treating a given condition. Bundled payments, by contrast, directly reward and motivate

cost reduction from the bottom up, team by team. At the same time, they encourage accurate cost measurement not only to inform price setting but to enable true cost reduction.

Bundled payments will be the catalyst that finally motivates provider teams to work together to understand the actual costs of each step in the entire care process, learn how to do things better, and get care right the first time. By encouraging competition for the treatment of individual conditions on the basis of quality and price, bundled payments also reward providers for standardizing care pathways, eliminating services and therapies that fail to improve outcomes, better utilizing staff to the top of their skills, and providing care in the right facilities. If providers use ineffective or unnecessary therapies or services, they will bear the cost, making bundled payments a check against overtreatment.

The result will be not just a downward “bend” in the cost curve—that is, a slower increase—but actual cost reduction. Our research suggests that savings of 20% to 30% are feasible in many conditions. And, because bundled payments are contingent on good outcomes, the right kind of cost reduction will take place, not cost cutting at the expense of quality.

Overcoming the Transition Challenges

Despite the now proven benefits of well-designed bundled payments, many hospital systems, group purchasing

organizations, private insurers, and some academics prefer capitation. Bundled payments, they argue, are too complicated to design, negotiate, and implement. (They ignore the fact that capitation models continue to rely on complex, expensive fee-for-service billing to pay clinicians and to set the baseline for calculating savings and penalties. Bundled payments are actually simpler to administer than the myriad of FFS payments for each patient over the care cycle.)

Skeptics raise a host of other objections: The scope of a condition and care cycle is hard to define; it is unrealistic to expect specialists to work together; the data on outcomes and costs needed to set prices are difficult to obtain; differences in risk across patients are hard to assess, which will lead to cherry-picking; and bundled payments won't rein in overtreatment.

A History of Success

BUNDLED PAYMENTS ARE NOT A NEW IDEA or a passing fad. Successful pilots date back for decades and include initiatives spearheaded by the Centers for Medicare & Medicaid Services.

Consider the Heart Bypass Demonstration, an initiative that ran from 1991 to 1996. CMS offered a bundled payment for coronary artery bypass graft surgery that covered all services delivered in the hospital, along with 90 days of post-discharge services. The pilot yielded savings to Medicare of $42.3 million, or roughly 10% of expected spending, at the

seven participating hospitals. The inpatient mortality rate declined at all the hospitals, and patient satisfaction improved.

CMS also implemented the Acute Care Episode program (from 2009 to 2011), in which Medicare paid five participating organizations a flat fee to cover hospital and physician services for various cardiac conditions and orthopedic care. Over a total of 12,501 episodes, the initiative generated an average savings to Medicare of 3.1% of expected costs.

If these objections represented serious barriers, we would expect to see little progress in implementing bundled payments and plenty of evidence that such programs were unsuccessful. To the contrary, bundled payments have a history of good results (see the sidebar "A History of Success") and are currently proliferating rapidly in a wide range of conditions, organizations, and countries.

In 2007, for example, the Netherlands introduced a successful bundled payment model for treating patients with type 2 diabetes, and, later, for chronic obstructive pulmonary disease (COPD). In 2009, the County of Stockholm, Sweden, introduced bundled payments for hip and knee replacements in healthy patients, achieving a 17% reduction in cost and a 33% reduction in complications over two years. More recently, Stockholm introduced bundled payments for all major spine diagnoses requiring surgery, and extensions to other conditions are under way there.

In 2011, Medicare introduced the voluntary Bundled Payments for Care Improvement (BPCI) program, which currently includes more than 14,000 bundles in 24 medical and 24 surgical conditions. Numerous physician practices have embraced the BPCI model, a transitional bundled payment approach that covers acute-care episodes and often a post-acute period of up to 90 days to promote better management of post-discharge services. According to participating providers, BPCI bundles have achieved significant improvements and savings an order of magnitude greater than savings from ACOs. Building on that success, CMS launched a mandatory bundled payment program for joint replacements in 2016, which covers 800 hospitals in 67 U.S. metropolitan areas.

Bundled payment contracts involving private insurers are also finally beginning to proliferate. For example, Twin Cities Orthopedics offers a bundle for joint replacement with most of the region's major insurers at a price well below the traditional hospital models. The practice reports better outcomes and cost reductions of more than 30%.

To be sure, many existing bundled payment programs have yet to encompass all the components of an ideal structure. Most have made pragmatic compromises, such as covering only part of the care cycle, using important but incomplete risk adjustments, and incorporating limited outcome measures. But even these less-than-comprehensive efforts are resulting in

major improvements, and the obstacles to bundled payments are being overcome.

Let's consider some of the main criticisms of bundled payments in more depth:

Only some conditions can be covered

Critics have suggested that bundled payments apply only to elective surgical care and other well-defined acute conditions, and not to nonsurgical conditions, chronic disease, or primary care. But this claim is inconsistent with actual experience. Of the 48 conditions designated for BPCI, only half were surgical. The other half were for care episodes in nonsurgical conditions, such as heart disease, kidney disease, diabetes, and COPD. Time-based bundled payments for chronic care are emerging in other countries and with private payers. Bundled payments work well for chronic conditions because of the huge benefits that result from coordinated longitudinal care by a multidisciplinary team.

Bundled payment models are also beginning to emerge for primary and preventive care for well-defined segments of patients with similar needs. Each primary care segment—such as healthy children, healthy adults, adults at risk for developing chronic disease, and the elderly—will need a very different mix of clinical, educational, and administrative services, and the appropriate outcomes will differ as well. Bundled payments reward integrated and efficient delivery of

the right mix of primary and preventive services for each patient group.

Primary care bundles need not cover the cost of treating complex, acute conditions, which are best paid for with bundled payments to IPUs covering those conditions. Instead, primary care teams should be held accountable for their performance in primary care and prevention for each patient segment: maintaining health status, avoiding disease progression, and preventing relapses.

Defining and implementing bundled payments is too complicated

Critics argue that it will be hard to negotiate bundled payments across all conditions and to get agreement on the definition of a medical condition, the extent of the care cycle, and the included services. This objection is weak at best. A manageable number of conditions account for a large proportion of health care costs, and we can start there and expand over time. The care required for most medical conditions is well established, and experience in defining bundles is rapidly accumulating. Methodologies and commercial tools, such as the use of comprehensive claims data sets, are in widespread use. Service companies that help providers define conditions, form teams, and manage payments are emerging, as are software tools that handle billing and claims processing for bundles.

Initially, bundled payments may cover less than the full care cycle, focus on simpler patient groups with a given condition,

and require adjudication mechanisms for gray areas that arise. This is already happening. As experience grows, bundled payments will become more comprehensive and inclusive. And a large body of evidence shows that the effort involved in understanding full care cycles and moving to multidisciplinary care is well worth it.

Why DRGs Are Not Bundled Payments

CRITICS OF BUNDLED PAYMENTS point to Medicare's experience with a superficially similar approach: the diagnosis-related group, or DRG, payment model. DRGs, which date back to 1984 and were adopted in many countries, were a step forward, but they did not trigger the hoped-for innovations in care delivery.

Why have DRGs failed to bring about greater change? DRGs make a single payment for a set of services provided at a given location; however, the payment does not cover the full care cycle for treating the patient's condition. By continuing to make separate payments to each specialist physician, hospital, and post-acute care site involved in a patient's care, DRGs perpetuate a system of uncoordinated care.

Moreover, DRG payments are not contingent on achieving good patient outcomes. Indeed, many DRGs fail to cover many support services crucial to good outcomes and overall value, such as patient education and counseling, behavioral health, and systematic follow-up. Under the DRG system, therefore, specialty silos in health care delivery have

remained largely intact. And providers continue to have no incentive to innovate to improve patient outcomes.

Providers won't work together

Critics argue that bundled payments hold providers accountable for care by other providers that they don't control; skeptics also claim that it will be hard to divide up a single payment to fairly recognize each party's contribution. This is one reason many hospital systems have been slow to embrace the new payment model. We are selling doctors short. Many physician groups have enthusiastically embraced bundles, because they see how the model rewards great care, motivates collaboration, and brings clinicians together. As physicians form condition-based IPUs and develop mechanisms for sharing accountability, formulas for dividing revenues and risk are emerging that reflect each provider's role, rather than flawed legacy fee structures.

At UCLA's kidney transplant program, for example, a bundled payment was first negotiated with several insurers more than 20 years ago. An IPU was formed and has become one of the premier U.S. kidney transplantation programs with superior outcomes. To divide the bundled price, urologists and nephrologists—the specialists who have the greatest impact on care—pay negotiated fees to other specialists involved in care (such as anesthesiology) and bear the residual financial risk

and share the gain. This structure has reinforced collaboration, not complicated it.

Another example is physician-owned OrthoCarolina's 2014 contract with Blue Cross and Blue Shield of North Carolina for bundled payment for joint replacement. OrthoCarolina provides care in several area hospitals and has negotiated a fixed payment with each of them for all the required inpatient care. Each participating hospital now has a designated team, including members of the nursing, quality, and administrative departments, that collaborates with OrthoCarolina surgeons in a virtual IPU. This ensures that everyone involved with the patient and the family fully understands the care pathway and expectations. The initial group of 220 patients in the plan experienced 0% readmissions, 0% reoperations, 0.45% deep venous thrombosis (versus 1% to 1.5% nationally), and substantial improvements in patient-reported quality-of-life outcomes. Average length of stay dropped from 2.4 days to 1.5 days, with 100% of patients discharged to their homes rather than a rehabilitation center. The cost per patient, as reported by Blue Cross and Blue Shield of North Carolina, fell an average of 20%.

Outcomes are difficult to measure

Critics claim that the outcome data at the medical condition level, an essential component of value-based bundled payments, doesn't exist or is too difficult and expensive to

collect. While this may have been true a decade ago, today outcome measurement is rapidly expanding, including patient-reported outcomes covering functional results crucial to patients. Many providers are already systematically measuring outcomes. Martini-Klinik, a high-volume IPU for prostate cancer in Hamburg, Germany, has been measuring a broad set of outcomes since its founding, in 1994. This has enabled it to achieve complication rates for impotence and incontinence that are far lower than average for Germany. In congenital heart disease care, Texas Children's tracks not only risk-adjusted surgical and intensive care mortality rates but also metrics of patients' neurodevelopmental status and, increasingly, ongoing quality of life.

Advances in information technology are making outcome measurement better, easier, less costly, and more reliable. Greater standardization of the set of outcomes to measure by condition will also make measurement more efficient and improve benchmarking. The International Consortium for Health Outcomes Measurement (ICHOM) has published global standard sets of outcomes and risk factors for 21 medical conditions that represent a significant portion of the disease burden, and the number is growing. Early bundled payment programs are already achieving significant outcome improvement. As provider experience grows, bundled payments will expand accountability and lead to even greater improvements.

Current cost information is inadequate

Critics argue that bundled payments require an understanding of costs that most providers lack, which puts them at unfair financial risk. Yet numerous bundled payment programs are already in place, using prices based on modest discounts from the sum of historical fee-for-service payments. New service companies are assisting providers in aggregating past charges and in reducing costs. Providers will learn to measure their actual costs, as organizations such as Mayo Clinic, MD Anderson, and the University of Utah are already doing. This will inform better price negotiations and accelerate cost reduction.

The failure of care delivery organizations to properly measure and manage costs is a crucial weakness in health care globally. Bundled payments will finally motivate providers to master proper costing and use cost data to drive efficiencies without sacrificing good patient outcomes.

Providers will cherry-pick patients

Critics charge that bundled payments will encourage providers to treat only the easiest and healthiest patients. But as we have already noted, proper bundled payments are risk-stratified or risk-adjusted. Even today's imperfect bundled payment contracts incorporate risk adjustments that are often better than those used in current FFS payment and beyond the crude risk adjustment used in capitation. Innovators are developing

pragmatic approaches that adjust for risk, such as restricting initial bundles to groups of patients with similar risk profiles for a condition. The County of Stockholm did this with joint replacements. Its initial bundle covered the 60% to 70% of patients classified as ASA 1 (normally healthy) or 2 (mild systemic disease); more-complex patients remained in the old reimbursement system. Careful tracking showed no evidence of bias in the selection of patients. The county plans to extend the bundle to more-complex joint replacement patients as better data becomes available.

Recently, the county introduced bundled payments for nine spine diagnoses requiring surgery, with far more sophisticated risk adjustment. The bundled payment includes a base payment, a payment covering expected complications, and a performance payment based on pain reduction. All three elements are adjusted for multiple patient risk factors. Risk adjustment will only improve as experience with it grows.

Bundled payments will encourage overtreatment

Critics raise concerns that bundled payments, like FFS, will lead to overtreatment because payment is tied to performing care, incenting providers to manufacture demand. Note that capitation plans, which have limited accountability for individual patient outcomes, have the opposite incentive: motivating providers to deny or delay the treatments patients need.

While definitive results are not yet available, our conversations with payers and government authorities in the United States, Sweden, and elsewhere have revealed no evidence that bundled payments have resulted in unnecessary surgeries or other treatments. Bundled payments are risk-adjusted and introduce transparency on outcomes, and the fixed payment will discourage unnecessary procedures, tests, and other services. Bundled payments (and all care) should incorporate appropriate use criteria (AUC), which use scientific evidence to define qualifications for particular treatments.

Price competition will trigger a race to the bottom

Finally, some providers worry that bundled payments will result in excessive price competition, as payers demand discounts and low-quality providers emerge offering cheap prices. This concern is common among hospitals, which are wary of greater competition and want to sustain existing reimbursement levels. We believe this fear is overblown. Bundled payments include clear accountability for outcomes and will penalize poor-quality providers. At the root of all these objections to bundled payments are critical failures that have held back health care for decades. Bundled payments will finally address these problems in ways that capitation cannot.

How Bundled Payments Will Transform Competition

As our multiple examples reveal, bundled payments are already transforming the way care is delivered. They unleash a new kind of competition that improves value for patients, informs and expands patient choice, lowers system cost, reshapes provider strategy, and alters industry structure for the better.

With bundled payments, patients are no longer locked into a single health system and can choose the provider that best meets their particular needs. Choice will expand dramatically as patients (and physicians) gain visibility into outcomes and prices of the providers that treat their condition. In a transparent bundled-payment world, patients will be able to decide whether to go to the hospital next door, travel across town, or venture even farther to a regional center of excellence for the care they need. This kind of choice, long overdue in health care, is what customers have in every other industry.

At the same time, the prices should fall. A bundled payment will usually be lower than the sum of current FFS reimbursements in today's inefficient and fragmented system. For conditions where legacy FFS payments failed to cover essential costs to achieve good outcomes, such as in mental health care or diagnostics that enable more targeted and successful treatments, prices may initially rise to support better care. But even these prices will fall as providers become more efficient.

In a world of bundled payments, market forces will determine provider prices and profitability, as they should. In today's system, FFS pricing allows inefficient or ineffective providers to be viable. With bundled payments, only providers that are effective and efficient will grow, earn attractive margins, and expand regionally and even nationally. The rest will see their margins decline, and those with poor outcomes will lose patients and bear the extra costs of dealing with avoidable complications, infections, readmissions, and repeat treatments.

Given today's hyperfragmentation of care, bundled payments should reduce the absolute number of providers treating each condition. But those that remain will be far stronger. And unlike the consolidation that would result from capitation, this winnowing of providers will create more-effective competition and greater accountability for results.

Providers will stop trying to do a little bit of everything and instead will target conditions where they can achieve good outcomes at low costs. Where they cannot, they will partner with more-effective providers or exit those service lines. The net result will be significantly better overall outcomes by condition and significantly lower average costs. No other payment model can produce such a transformation.

The shift to bundled payments will also spill over to drive positive change in pharmaceuticals, medical devices, diagnostic testing, imaging, and other suppliers. Today,

suppliers compete to get on approved lists, curry favor with prescribing specialists through consulting and research payments, and advertise directly to patients so that they will ask their doctor for particular treatments. As a result, many patients receive therapies that are not the best option, deliver little benefit, or are unnecessary. With bundled payments, suppliers will have to demonstrate that their particular drug, device, diagnostic test, or imaging method actually improves outcomes, lowers the overall cost, or both. Suppliers that can demonstrate value will command fair prices and gain market share, and there will be substantial cost reduction in the system overall. Competition on value is the best way to control the costs of expensive drugs and therapies, not today's approach of restricting access or attacking high prices as unethical or evil regardless of the value products offer.

The Time Is Now

The biggest beneficiary of bundled payments will be patients, who will receive better care and have access to more choice. The best providers will also prosper. Many already recognize that bundled payments enable them to compete on value, transform care, and put the system on a sustainable health care path for the long run. Those already organized into IPUs for specific medical conditions are particularly well-positioned to

move aggressively. Physician groups in particular have often moved the fastest.

Many health systems, however, have been reluctant to get behind bundled payments. They seem to believe that capitation better preserves the status quo—a top-down approach that leverages their clout and scale. They also see it as encouraging industry consolidation, which will ease reimbursement pressure and reduce competition. However, leading health systems are embracing bundled payments and the shift in competition to what really matters to patients.

Health systems with their own insurance plans, or those that self-insure care for their employees, can begin immediately to introduce bundled payments internally. Health systems that have adopted ACOs or other capitated models can also use condition-based bundled payments to pay internal units. Doing so will accelerate learning while motivating clinical units to improve outcomes and reduce costs in a way that existing departmental budgets or FFS can never match. Adopting bundles internally will be a stepping stone to contracting this way with payers and directly with employers.

Payers will reap huge benefits from bundled payments. Single-payer systems, such as those in Canada, Sweden, and the U.S. Veterans Administration, are well-positioned to transition to bundled payments for a growing number of medical conditions. Indeed, this is already happening in some

countries and regions, with CMS leading the way in the United States.

But many private insurers, which have prospered under the status quo, have been disappointingly slow in moving to bundled payments. Many seem to favor capitation as less of a change; they believe it preserves payment infrastructure while shifting risk to providers. As an excuse, they cite their inability to process claims for bundled payments, even though bundled claims processing is inherently far simpler.

Improving the way they pay for health care, however, is the only means by which insurers can offer greater value to its customers. Insurers must do so, or they will have a diminished role in the system. We challenge the industry to shift from being the obstacle to bundled payment to becoming the driver. Recently, we've been heartened to see more private insurers moving toward bundled payments.

Employers, which actually pay for much of health insurance in the United States, should step up to lead the move to bundled payments. This will improve outcomes for their employees, bring down prices, and increase competition. Self-insured employer health plans need to direct their plan administrators to roll out bundles, starting with costly conditions for which employees experience uneven outcomes.

Should their insurers fail to move toward bundles, large employers have the clout to go directly to providers. Lowe's, Boeing, and Walmart are contracting directly with providers

such as Mayo Clinic, Cleveland Clinic, Virginia Mason, and Geisinger on bundled payments for orthopedics and complex cardiac care. The Health Transformation Alliance, consisting of 20 large employers that account for 4 million lives, is pooling data and purchasing power to accelerate the implementation of bundled payments.

The time has come to change the way we pay for health care, in the United States and around the world. Capitation is not the solution. It entrenches large existing systems, eliminates patient choice, promotes more consolidation, limits competition, and perpetuates the lack of provider accountability for outcomes. It will fail again to drive true innovation in health care delivery.

Capitation will also fail to stem the tide of the ever-rising costs of health care. ACOs, despite their strong advocates, have produced minimal cost savings (0.1%). By contrast, even the simplified bundled payment contracts under way today are achieving better results. Medicare is expected to save at least 2% ($250 million) in its program's first full year of operation. And experience in the United States and elsewhere shows that the savings can be far larger.

Capitation might seem simple, but given highly heterogeneous populations and continual turnover of patients and physicians, it is actually harder to implement, risk-adjust,

and manage to deliver improved care. Bundled payments, in contrast, are a direct and intuitive way to pay clinical teams for delivering value, condition by condition. They put accountability where it should be—on outcomes that matter to patients. This way to pay for health care is working, and expanding rapidly.

Much remains to be done to put bundled payments into widespread practice, but the barriers are rapidly being overcome. Bundled payments are the only true value-based payment model for health care. The time is now.

Further Reading

- **"The Strategy That Will Fix Health Care,"** Michael E. Porter and Thomas H. Lee, HBR, October 2013
- **"What Is Value in Health Care?,"** Michael E. Porter, *New England Journal of Medicine,* December 2010
- **"How to Solve the Cost Crisis in Health Care,"** Robert S. Kaplan and Michael E. Porter, HBR, September 2011
- **"Redesigning Primary Care: A Strategic Vision to Improve Value by Organizing Around Patients' Needs,"** Michael E. Porter, Erika A. Pabo, and Thomas H. Lee, *Health Affairs,* March 2013
- **"Getting Bundled Payments Right in Health Care,"** Derek A. Haas, Robert S. Kaplan, Dereesa Reid, Jonathan Warsh, and Michael E.

West, HBR.org, October 2015

Originally published in July–August 2016. Reprint R1607G

The Performance Management Revolution

by Peter Cappelli and Anna Tavis

WHEN BRIAN JENSEN TOLD HIS AUDIENCE of HR executives that Colorcon wasn't bothering with annual reviews anymore, they were appalled. This was in 2002, during his tenure as the drugmaker's head of global human resources. In his presentation at the Wharton School, Jensen explained that Colorcon had found a more effective way of reinforcing desired behaviors and managing performance: Supervisors were giving people instant feedback, tying it to individuals' own goals, and handing out small weekly bonuses to employees they saw doing good things.

Back then the idea of abandoning the traditional appraisal process—and all that followed from it—seemed heretical. But now, by some estimates, more than one-third of U.S. companies are doing just that. From Silicon Valley to New York, and in offices across the world, firms are replacing annual reviews

with frequent, informal check-ins between managers and employees.

As you might expect, technology companies such as Adobe, Juniper Systems, Dell, Microsoft, and IBM have led the way. Yet they've been joined by a number of professional services firms (Deloitte, Accenture, PwC), early adopters in other industries (Gap, Lear, OppenheimerFunds), and even General Electric, the longtime role model for traditional appraisals.

Without question, rethinking performance management is at the top of many executive teams' agendas, but what drove the change in this direction? Many factors. In a recent article for *People + Strategy,* a Deloitte manager referred to the review process as "an investment of 1.8 million hours across the firm that didn't fit our business needs anymore." One *Washington Post* business writer called it a "rite of corporate kabuki" that restricts creativity, generates mountains of paperwork, and serves no real purpose. Others have described annual reviews as a last-century practice and blamed them for a lack of collaboration and innovation. Employers are also finally acknowledging that both supervisors and subordinates despise the appraisal process—a perennial problem that feels more urgent now that the labor market is picking up and concerns about retention have returned.

But the biggest limitation of annual reviews—and, we have observed, the main reason more and more companies are dropping them—is this: With their heavy emphasis on financial

rewards and punishments and their end-of-year structure, they hold people accountable for past behavior at the expense of improving current performance and grooming talent for the future, both of which are critical for organizations' long-term survival. In contrast, regular conversations about performance and development change the focus to building the workforce your organization needs to be competitive both today and years from now. Business researcher Josh Bersin estimates that about 70% of multinational companies are moving toward this model, even if they haven't arrived quite yet.

The tension between the traditional and newer approaches stems from a long-running dispute about managing people: Do you "get what you get" when you hire your employees? Should you focus mainly on motivating the strong ones with money and getting rid of the weak ones? Or are employees malleable? Can you change the way they perform through effective coaching and management and intrinsic rewards such as personal growth and a sense of progress on the job?

With traditional appraisals, the pendulum had swung too far toward the former, more transactional view of performance, which became hard to support in an era of low inflation and tiny merit-pay budgets. Those who still hold that view are railing against the recent emphasis on improvement and growth over accountability. But the new perspective is unlikely to be a flash in the pan because, as we will discuss, it is being driven by business needs, not imposed by HR.

Idea in Brief

The Problem

By emphasizing individual accountability for past results, traditional appraisals give short shrift to improving current performance and developing talent for the future. That can hinder long-term competitiveness.

The Solution

To better support employee development, many organizations are dropping or radically changing their annual review systems in favor of giving people less formal, more frequent feedback that follows the natural cycle of work.

The Outlook

This shift isn't just a fad—real business needs are driving it. Support at the top is critical, though. Some firms that have struggled to go entirely without ratings are trying a "third way": assigning multiple ratings several times a year to encourage employees' growth.

First, though, let's consider how we got to this point—and how companies are faring with new approaches.

How We Got Here

Historical and economic context has played a large role in the evolution of performance management over the decades. When human capital was plentiful, the focus was on which

people to let go, which to keep, and which to reward—and for those purposes, traditional appraisals (with their emphasis on individual accountability) worked pretty well. But when talent was in shorter supply, as it is now, developing people became a greater concern—and organizations had to find new ways of meeting that need.

From accountability to development

Appraisals can be traced back to the U.S. military's "merit rating" system, created during World War I to identify poor performers for discharge or transfer. After World War II, about 60% of U.S. companies were using them (by the 1960s, it was closer to 90%). Though seniority rules determined pay increases and promotions for unionized workers, strong merit scores meant good advancement prospects for managers. At least initially, *improving* performance was an afterthought.

And then a severe shortage of managerial talent caused a shift in organizational priorities: Companies began using appraisals to develop employees into supervisors, and especially managers into executives. In a famous 1957 HBR article, social psychologist Douglas McGregor argued that subordinates should, with feedback from the boss, help set their performance goals and assess themselves—a process that would build on their strengths and potential. This "Theory Y" approach to management—he coined the term later on—assumed that employees wanted to perform well and would do

so if supported properly. ("Theory X" assumed you had to motivate people with material rewards and punishments.) McGregor noted one drawback to the approach he advocated: Doing it right would take managers several days per subordinate each year.

By the early 1960s, organizations had become so focused on developing future talent that many observers thought that tracking past performance had fallen by the wayside. Part of the problem was that supervisors were reluctant to distinguish good performers from bad. One study, for example, found that 98% of federal government employees received "satisfactory" ratings, while only 2% got either of the other two outcomes: "unsatisfactory" or "outstanding." After running a well-publicized experiment in 1964, General Electric concluded it was best to split the appraisal process into separate discussions about accountability and development, given the conflicts between them. Other companies followed suit.

Back to accountability

In the 1970s, however, a shift began. Inflation rates shot up, and merit-based pay took center stage in the appraisal process. During that period, annual wage increases really mattered. Supervisors often had discretion to give raises of 20% or more to strong performers, to distinguish them from the sea of employees receiving basic cost-of-living raises, and getting no increase represented a substantial pay cut. With the stakes so

high—and with antidiscrimination laws so recently on the books—the pressure was on to award pay more objectively. As a result, accountability became a higher priority than development for many organizations.

Three other changes in the zeitgeist reinforced that shift:

First, Jack Welch became CEO of General Electric in 1981. To deal with the long-standing concern that supervisors failed to label real differences in performance, Welch championed the forced-ranking system—another military creation. Though the U.S. Army had devised it, just before entering World War II, to quickly identify a large number of officer candidates for the country's imminent military expansion, GE used it to shed people at the bottom. Equating performance with individuals' inherent capabilities (and largely ignoring their potential to grow), Welch divided his workforce into "A" players, who must be rewarded; "B" players, who should be accommodated; and "C" players, who should be dismissed. In that system, development was reserved for the "A" players—the high-potentials chosen to advance into senior positions.

Second, 1993 legislation limited the tax deductibility of executive salaries to $1 million but exempted performance-based pay. That led to a rise in outcome-based bonuses for corporate leaders—a change that trickled down to frontline managers and even hourly employees—and organizations relied even more on the appraisal process to assess merit.

Third, McKinsey's War for Talent research project in the late 1990s suggested that some employees were fundamentally more talented than others (you knew them when you saw them, the thinking went). Because such individuals were, by definition, in short supply, organizations felt they needed to take great care in tracking and rewarding them. Nothing in the McKinsey studies showed that fixed personality traits actually made certain people perform better, but that was the assumption.

So, by the early 2000s, organizations were using performance appraisals mainly to hold employees accountable and to allocate rewards. By some estimates, as many as one-third of U.S. corporations—and 60% of the *Fortune* 500—had adopted a forced-ranking system. At the same time, other changes in corporate life made it harder for the appraisal process to advance the time-consuming goals of improving individual performance and developing skills for future roles. Organizations got much flatter, which dramatically increased the number of subordinates that supervisors had to manage. The new norm was 15 to 25 direct reports (up from six before the 1960s). While overseeing more employees, supervisors were also expected to be individual contributors. So taking days to manage the performance issues of each employee, as Douglas McGregor had advocated, was impossible. Meanwhile, greater interest in lateral hiring reduced the need for internal

development. Up to two-thirds of corporate jobs were filled from outside, compared with about 10% a generation earlier.

Back to development … again

Another major turning point came in 2005: A few years after Jack Welch left GE, the company quietly backed away from forced ranking because it fostered internal competition and undermined collaboration. Welch still defends the practice, but what he really supports is the general principle of letting people know how they are doing: "As a manager, you owe candor to your people," he wrote in the *Wall Street Journal* in 2013. "They must not be guessing about what the organization thinks of them." It's hard to argue against candor, of course. But more and more firms began questioning how useful it was to compare people with one another or even to rate them on a scale.

So the emphasis on accountability for past performance started to fade. That continued as jobs became more complex and rapidly changed shape—in that climate, it was difficult to set annual goals that would still be meaningful 12 months later. Plus, the move toward team-based work often conflicted with individual appraisals and rewards. And low inflation and small budgets for wage increases made appraisal-driven merit pay seem futile. What was the point of trying to draw performance distinctions when rewards were so trivial?

The whole appraisal process was loathed by employees anyway. Social science research showed that they hated numerical scores—they would rather be told they were “average” than given a 3 on a 5-point scale. They especially detested forced ranking. As Wharton’s Iwan Barankay demonstrated in a field setting, performance actually declined when people were rated relative to others. Nor did the ratings seem accurate. As the accumulating research on appraisal scores showed, they had as much to do with who the rater was (people gave higher ratings to those who were like them) as they did with performance.

And managers hated *doing* reviews, as survey after survey made clear. Willis Towers Watson found that 45% did not see value in the systems they used. Deloitte reported that 58% of HR executives considered reviews an ineffective use of supervisors’ time. In a study by the advisory service CEB, the average manager reported spending about 210 hours—close to five weeks—doing appraisals each year.

As dissatisfaction with the traditional process mounted, high-tech firms ushered in a new way of thinking about performance. The “Agile Manifesto,” created by software developers in 2001, outlined several key values—favoring, for instance, “responding to change over following a plan.” It emphasized principles such as collaboration, self-organization, self-direction, and regular reflection on how to work more effectively, with the aim of prototyping more quickly and

responding in real time to customer feedback and changes in requirements. Although not directed at performance per se, these principles changed the definition of effectiveness on the job—and they were at odds with the usual practice of cascading goals from the top down and assessing people against them once a year.

So it makes sense that the first significant departure from traditional reviews happened at Adobe, in 2011. The company was already using the agile method, breaking down projects into “sprints” that were immediately followed by debriefing sessions. Adobe explicitly brought this notion of constant assessment and feedback into performance management, with frequent check-ins replacing annual appraisals. Juniper Systems, Dell, and Microsoft were prominent followers.

CEB estimated in 2014 that 12% of U.S. companies had dropped annual reviews altogether. Willis Towers Watson put the figure at 8% but added that 29% were considering eliminating them or planning to do so. Deloitte reported in 2015 that only 12% of the U.S. companies it surveyed were *not* planning to rethink their performance management systems. This trend seems to be extending beyond the United States as well. PwC reports that two-thirds of large companies in the UK, for example, are in the process of changing their systems.

Three Business Reasons to Drop Appraisals

In light of that history, we see three clear business imperatives that are leading companies to abandon performance appraisals:

The return of people development

Companies are under competitive pressure to upgrade their talent management efforts. This is especially true at consulting and other professional services firms, where knowledge work is the offering—and where inexperienced college grads are turned into skilled advisers through structured training. Such firms are doubling down on development, often by putting their employees (who are deeply motivated by the potential for learning and advancement) in charge of their own growth. This approach requires rich feedback from supervisors—a need that's better met by frequent, informal check-ins than by annual reviews.

Now that the labor market has tightened and keeping good people is once again critical, such companies have been trying to eliminate "dissatisfiers" that drive employees away. Naturally, annual reviews are on that list, since the process is so widely reviled and the focus on numerical ratings interferes with the learning that people want and need to do. Replacing this system with feedback that's delivered right after client engagements helps managers do a better job of coaching and allows subordinates to process and apply the advice more effectively.

Kelly Services was the first big professional services firm to drop appraisals, in 2011. PwC tried it with a pilot group in 2013 and then discontinued annual reviews for all 200,000-plus employees. Deloitte followed in 2015, and Accenture and KPMG made similar announcements shortly thereafter. Given the sheer size of these firms, and the fact that they offer management advice to thousands of organizations, their choices are having an enormous impact on other companies. Firms that scrap appraisals are also rethinking employee management much more broadly. Accenture CEO Pierre Nanterme estimates that his firm is changing about 90% of its talent practices.

The need for agility

When rapid innovation is a source of competitive advantage, as it is now in many companies and industries, that means future needs are continually changing. Because organizations won't necessarily want employees to keep doing the same things, it doesn't make sense to hang on to a system that's built mainly to assess and hold people accountable for past or current practices. As Susan Peters, GE's head of human resources, has pointed out, businesses no longer have clear annual cycles. Projects are short-term and tend to change along the way, so employees' goals and tasks can't be plotted out a year in advance with much accuracy.

At GE a new business strategy based on innovation was the biggest reason the company recently began eliminating individual ratings and annual reviews. Its new approach to performance management is aligned with its FastWorks platform for creating products and bringing them to market, which borrows a lot from agile techniques. Supervisors still have an end-of-year summary discussion with subordinates, but the goal is to push frequent conversations with employees (GE calls them "touchpoints") and keep revisiting two basic questions: What am I doing that I should keep doing? And what am I doing that I should change? Annual goals have been replaced with shorter-term "priorities." As with many of the companies we see, GE first launched a pilot, with about 87,000 employees in 2015, before adopting the changes across the company.

The centrality of teamwork

Moving away from forced ranking and from appraisals' focus on individual accountability makes it easier to foster teamwork. This has become especially clear at retail companies like Sears and Gap—perhaps the most surprising early innovators in appraisals. Sophisticated customer service now requires frontline and back-office employees to work together to keep shelves stocked and manage customer flow, and traditional systems don't enhance performance at the team level or help track collaboration.

Gap supervisors still give workers end-of-year assessments, but only to summarize performance discussions that happen throughout the year and to set pay increases accordingly. Employees still have goals, but as at other companies, the goals are short-term (in this case, quarterly). Now two years into its new system, Gap reports far more satisfaction with its performance process and the best-ever completion of store-level goals. Nonetheless, Rob Ollander-Krane, Gap's senior director of organization performance effectiveness, says the company needs further improvement in setting stretch goals and focusing on team performance.

Implications

All three reasons for dropping annual appraisals argue for a system that more closely follows the natural cycle of work. Ideally, conversations between managers and employees occur when projects finish, milestones are reached, challenges pop up, and so forth—allowing people to solve problems in current performance while also developing skills for the future. At most companies, managers take the lead in setting near-term goals, and employees drive career conversations throughout the year. In the words of one Deloitte manager: "The conversations are more holistic. They're about goals and strengths, not just about past performance."

Perhaps most important, companies are overhauling performance management because their businesses require the

change. That's true whether they're professional services firms that must develop people in order to compete, companies that need to deliver ongoing performance feedback to support rapid innovation, or retailers that need better coordination between the sales floor and the back office to serve their customers.

Of course, many HR managers worry: If we can't get supervisors to have good conversations with subordinates once a year, how can we expect them to do so more frequently, without the support of the usual appraisal process? It's a valid question—but we see reasons to be optimistic.

As GE found in 1964 and as research has documented since, it is extraordinarily difficult to have a serious, open discussion about problems while also dishing out consequences such as low merit pay. The end-of-year review was also an excuse for delaying feedback until then, at which point both the supervisor and the employee were likely to have forgotten what had happened months earlier. Both of those constraints disappear when you take away the annual review. Additionally, almost all companies that have dropped traditional appraisals have invested in training supervisors to talk more about development with their employees—and they are checking with subordinates to make sure that's happening.

Moving to an informal system requires a culture that will keep the continuous feedback going. As Megan Taylor, Adobe's director of business partnering, pointed out at a recent conference, it's difficult to sustain that if it's not happening

organically. Adobe, which has gone totally numberless but still gives merit increases based on informal assessments, reports that regular conversations between managers and their employees are now occurring without HR's prompting. Deloitte, too, has found that its new model of frequent, informal check-ins has led to more meaningful discussions, deeper insights, and greater employee satisfaction. (For more details, see "Reinventing Performance Management," HBR, April 2015.) The firm started to go numberless like Adobe but then switched to assigning employees several numbers four times a year, to give them rolling feedback on different dimensions. Jeffrey Orlando, who heads up development and performance at Deloitte, says the company has been tracking the effects on business results, and they've been positive so far.

Challenges That Persist

The greatest resistance to abandoning appraisals, which is something of a revolution in human resources, comes from HR itself. The reason is simple: Many of the processes and systems that HR has built over the years revolve around those performance ratings. Experts in employment law had advised organizations to standardize practices, develop objective criteria to justify every employment decision, and document all relevant facts. Taking away appraisals flies in the face of that

advice—and it doesn't necessarily solve every problem that they failed to address.

A talent management timeline

The tug-of-war between accountability and development over the decades

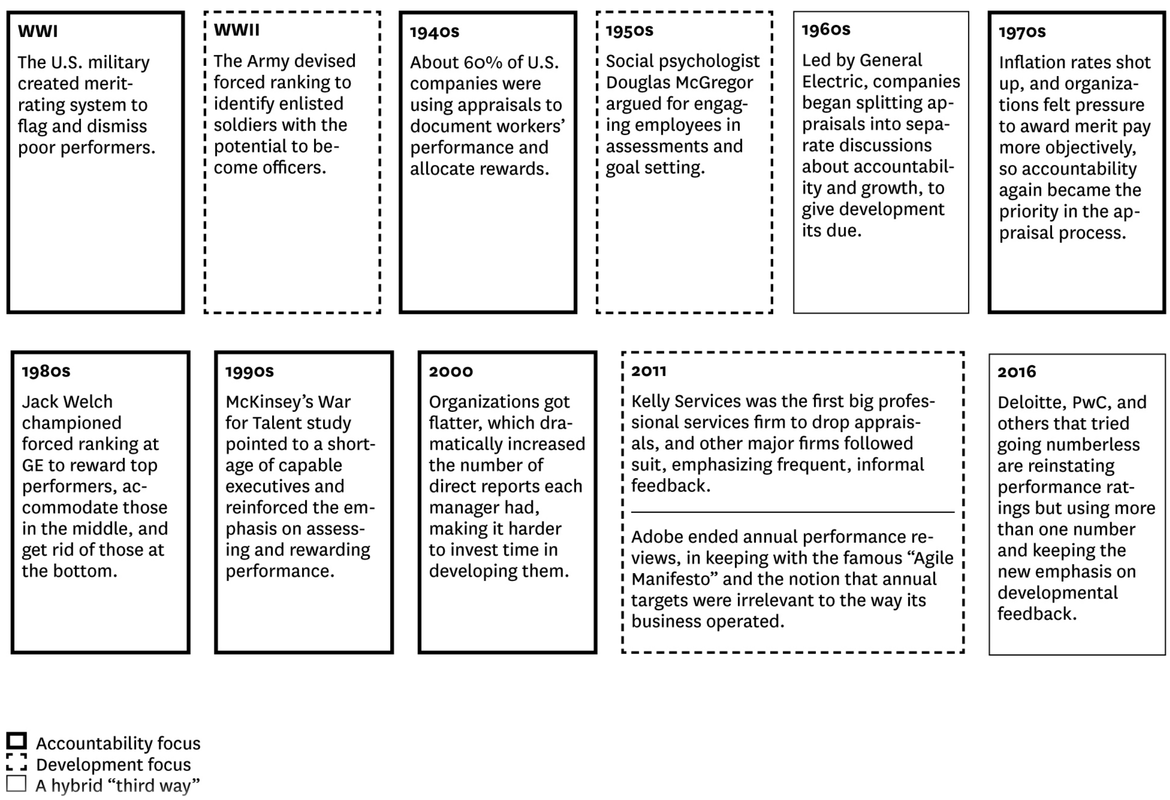

Here are some of the challenges that organizations still grapple with when they replace the old performance model with new approaches:

Aligning individual and company goals

In the traditional model, business objectives and strategies cascaded down the organization. All the units, and then all the individual employees, were supposed to establish their goals to reflect and reinforce the direction set at the top. But this approach works only when business goals are easy to articulate and held constant over the course of a year. As we've discussed, that's often not the case these days, and employee goals may be pegged to specific projects. So as projects unfold and tasks change, how do you coordinate individual priorities with the goals for the whole enterprise, especially when the business objectives are short-term and must rapidly adapt to market shifts? It's a new kind of problem to solve, and the jury is still out on how to respond.

Rewarding performance

Appraisals gave managers a clear-cut way of tying rewards to individual contributions. Companies changing their systems are trying to figure out how their new practices will affect the pay-for-performance model, which none of them have explicitly abandoned.

They still differentiate rewards, usually relying on managers' qualitative judgments rather than numerical ratings. In pilot programs at Juniper Systems and Cargill, supervisors had no difficulty allocating merit-based pay without appraisal scores. In fact, both line managers and HR staff felt that paying closer

attention to employee performance throughout the year was likely to make their merit-pay decisions more valid.

But it will be interesting to see whether most supervisors end up reviewing the feedback they've given each employee over the year before determining merit increases. (Deloitte's managers already do this.) If so, might they produce something *like* an annual appraisal score—even though it's more carefully considered? And could that subtly undermine development by shifting managers' focus back to accountability?

Identifying poor performers

Though managers may assume they need appraisals to determine which employees aren't doing their jobs well, the traditional process doesn't *really* help much with that. For starters, individuals' ratings jump around over time. Research shows that last year's performance score predicts only one-third of the variance in this year's score—so it's hard to say that someone simply isn't up to scratch. Plus, HR departments consistently complain that line managers don't use the appraisal process to document poor performers. Even when they do, waiting until the end of the year to flag struggling employees allows failure to go on for too long without intervention.

We've observed that companies that have dropped appraisals are requiring supervisors to immediately identify problem employees. Juniper Systems also formally asks supervisors each

quarter to confirm that their subordinates are performing up to company standards. Only 3%, on average, are not, and HR is brought in to address them. Adobe reports that its new system has reduced dismissals, because struggling employees are monitored and coached much more closely.

Still, given how reluctant most managers are to single out failing employees, we can't assume that getting rid of appraisals will make those tough calls any easier. And all the companies we've observed still have "performance improvement plans" for employees identified as needing support. Such plans remain universally problematic, too, partly because many issues that cause poor performance can't be solved by management intervention.

Avoiding legal troubles

Employee relations managers within HR often worry that discrimination charges will spike if their companies stop basing pay increases and promotions on numerical ratings, which seem objective. But appraisals haven't prevented discriminatory practices. Though they force managers to systematically review people's contributions each year, a great deal of discretion (always subject to bias) is built into the process, and considerable evidence shows that supervisors discriminate against some employees by giving them undeservedly low ratings.

Leaders at Gap report that their new practices were driven partly by complaints and research showing that the appraisal process was often biased and ineffective. Frontline workers in retail (disproportionately women and minorities) are especially vulnerable to unfair treatment. Indeed, formal ratings may do more to *reveal* bias than to curb it. If a company has clear appraisal scores and merit-pay indexes, it is easy to see if women and minorities with the same scores as white men are getting fewer or lower pay increases.

All that said, it's not clear that new approaches to performance management will do much to mitigate discrimination either. (See the sidebar "Can You Take Cognitive Bias Out of Assessments?") Gap has found that getting rid of performance scores increased fairness in pay and other decisions, but judgments still have to be made—and there's the possibility of bias in every piece of qualitative information that decision makers consider.

Managing the feedback firehose

In recent years most HR information systems were built to move annual appraisals online and connect them to pay increases, succession planning, and so forth. They weren't designed to accommodate continuous feedback, which is one reason many employee check-ins consist of oral comments, with no documentation.

The tech world has responded with apps that enable supervisors to give feedback anytime and to record it if desired. At General Electric, the PD@GE app (“PD” stands for “performance development”) allows managers to call up notes and materials from prior conversations and summarize that information. Employees can use the app to ask for direction when they need it. IBM has a similar app that adds another feature: It enables employees to give feedback to peers and choose whether the recipient’s boss gets a copy. Amazon’s Anytime Feedback tool does much the same thing. The great advantage of these apps is that supervisors can easily review all the discussion text when it is time to take actions such as award merit pay or consider promotions and job reassignments.

Can You Take Cognitive Bias out of Assessments?

A CLASSIC STUDY BY EDWARD JONES and Victor Harris in the 1960s demonstrated that people tend to attribute others’ behavior to character rather than circumstances.

When a car goes streaking past us, for instance, we think that the driver is a jerk and ignore the possibility that there might be an emergency. A good workplace example of this cognitive bias—known as the “fundamental attribution error”—is to assume that the lowest performers in any year will always be

the worst performers and to fire them as a result. Such an assumption overlooks the impact of good or poor management, not to mention business conditions that are beyond employees' control.

Of course, this model is highly flattering to people who have advanced into executive roles—"A" players whose success is, by definition, credited to their superior abilities, not to good fortune. That may be partly why the model has persisted so long in the face of considerable evidence against it.

Even when "A" players seem to perform well in many contexts (and that's rarely measured), they may be coasting on the "halo effect"—another type of bias, akin to self-fulfilling prophecy. If these folks have already been successful, they receive more opportunities than others, and they're pushed harder, so naturally they do better.

Biases color individual performance ratings as well. Decision makers may give past behavior too much weight, for instance, or fall prey to stereotypes when they assign their ratings.

But when you get rid of forced ranking and appraisal scores, you don't eradicate bias. Discrimination and faulty assumptions still creep into qualitative assessments. In some ways the older, more cumbersome performance systems actually made it harder for managers to keep their blinders on. Formal feedback from various stakeholders provided some balance when supervisors were otherwise inclined to see only the good things their stars did and failed to recognize others' contributions.

Anytime you exercise judgment, whether or not you translate that to numerical ratings, intuition plays a part, and bias can rear its head.

Of course, being on the receiving end of all that continual coaching could get overwhelming—it never lets up. And as for peer feedback, it isn't always useful, even if apps make it easier to deliver in real time. Typically, it's less objective than supervisor feedback, as anyone familiar with 360s knows. It can be also "gamed" by employees to help or hurt colleagues. (At Amazon, the cutthroat culture encourages employees to be critical of one another's performance, and forced ranking creates an incentive to push others to the bottom of the heap.) The more consequential the peer feedback, the more likely the problems.

Not all employers face the same business pressures to change their performance processes. In some fields and industries (think sales and financial services), it still makes sense to emphasize accountability and financial rewards for individual performers. Organizations with a strong public mission may also be well served by traditional appraisals. But even government organizations like NASA and the FBI are rethinking their approach, having concluded that accountability should be collective and that supervisors need to do a better job of coaching and developing their subordinates.

Ideology at the top matters. Consider what happened at Intel. In a two-year pilot, employees got feedback but no formal appraisal scores. Though supervisors did not have difficulty differentiating performance or distributing performance-based pay without the ratings, company executives returned to using them, believing they created healthy competition and clear outcomes. At Sun Communities, a manufactured-home company, senior leaders also oppose eliminating appraisals because they think formal feedback is essential to accountability. And Medtronic, which gave up ratings several years ago, is resurrecting them now that it has acquired Ireland-based Covidien, which has a more traditional view of performance management.

Other firms aren't completely reverting to old approaches but instead seem to be seeking middle ground. As we've mentioned, Deloitte has backpedaled from giving no ratings at all to having project leads and managers assign them in four categories on a quarterly basis, to provide detailed "performance snapshots." PwC recently made a similar move in its client-services practices: Employees still don't receive a single rating each year, but they now get scores on five competencies, along with other development feedback. In PwC's case, the pushback against going numberless actually came from employees, especially those on a partner track, who wanted to know how they were doing.

At one insurance company, after formal ratings had been eliminated, merit-pay increases were being shared internally and then interpreted as performance scores. These became known as “shadow ratings,” and because they started to affect other talent management decisions, the company eventually went back to formal appraisals. But it kept other changes it had made to its performance management system, such as quarterly conversations between managers and employees, to maintain its new commitment to development.

It will be interesting to see how well these “third way” approaches work. They, too, could fail if they aren’t supported by senior leadership and reinforced by organizational culture. Still, in most cases, sticking with old systems seems like a bad option. Companies that don’t think an overhaul makes sense for them should at least carefully consider whether their process is giving them what they need to solve current performance problems and develop future talent. Performance appraisals wouldn’t be the least popular practice in business, as they’re widely believed to be, if *something* weren’t fundamentally wrong with them.

Originally published in October 2016. Reprint R1610D

Let Your Workers Rebel

by Francesca Gino

THROUGHOUT OUR CAREERS, we are taught to conform—to the status quo, to the opinions and behaviors of others, and to information that supports our views. The pressure only grows as we climb the organizational ladder. By the time we reach high-level positions, conformity has been so hammered into us that we perpetuate it in our enterprises. In a recent survey I conducted of more than 2,000 employees across a wide range of industries, nearly half the respondents reported working in organizations where they regularly feel the need to conform, and more than half said that people in their organizations do not question the status quo. The results were similar when I surveyed high-level executives and midlevel managers. As this data suggests, organizations consciously or unconsciously urge employees to check a good chunk of their real selves at the door. Workers and their organizations both pay a price: decreased engagement, productivity, and innovation (see the exhibit "The perils of conformity").

Drawing on my research and fieldwork and on the work of other scholars of psychology and management, I will describe three reasons for our conformity on the job, discuss why this behavior is costly for organizations, and suggest ways to combat it.

Of course, not all conformity is bad. But to be successful and evolve, organizations need to strike a balance between adherence to the formal and informal rules that provide necessary structure and the freedom that helps employees do their best work. The pendulum has swung too far in the direction of conformity. In another recent survey I conducted, involving more than 1,000 employees in a variety of industries, less than 10% said they worked in companies that regularly encourage nonconformity. That's not surprising: For decades the principles of scientific management have prevailed. Leaders have been overly focused on designing efficient processes and getting employees to follow them. Now they need to think about when conformity hurts their business and allow—even promote—what I call *constructive nonconformity:* behavior that deviates from organizational norms, others' actions, or common expectations, to the benefit of the organization.

Why Conformity Is So Prevalent

Let's look at the three main, and interrelated, reasons why we so often conform at work.

We fall prey to social pressure

Early in life we learn that tangible benefits arise from following social rules about what to say, how to act, how to dress, and so on. Conforming makes us feel accepted and part of the majority. As classic research conducted in the 1950s by the psychologist Solomon Asch showed, conformity to peer pressure is so powerful that it occurs even when we know it will lead us to make bad decisions. In one experiment, Asch asked participants to complete what they believed was a simple perceptual task: identifying which of three lines on one card was the same length as a line on another card. When asked individually, participants chose the correct line. When asked in the presence of paid actors who intentionally selected the wrong line, about 75% conformed to the group at least once. In other words, they chose an incorrect answer in order to fit in.

Organizations have long exploited this tendency. Ancient Roman families employed professional mourners at funerals. Entertainment companies hire people ("claques") to applaud at performances. And companies advertising health products often report the percentage of doctors or dentists who use their offerings.

Conformity at work takes many forms: modeling the behavior of others in similar roles, expressing appropriate emotions, wearing proper attire, routinely agreeing with the opinions of managers, acquiescing to a team's poor decisions, and so on. And all too often, bowing to peer pressure reduces

individuals' engagement with their jobs. This is understandable: Conforming often conflicts with our true preferences and beliefs and therefore makes us feel inauthentic. In fact, research I conducted with Maryam Kouchaki, of Northwestern University, and Adam Galinsky, of Columbia University, showed that when people feel inauthentic at work, it's usually because they have succumbed to social pressure to conform.

We become too comfortable with the status quo

In organizations, standard practices—the usual ways of thinking and doing—play a critical role in shaping performance over time. But they can also get us stuck, decrease our engagement, and constrain our ability to innovate or to perform at a high level. Rather than resulting from thoughtful choices, many traditions endure out of routine, or what psychologists call the *status quo bias*. Because we feel validated and reassured when we stick to our usual ways of thinking and doing, and because—as research has consistently found—we weight the potential losses of deviating from the status quo much more heavily than we do the potential gains, we favor decisions that maintain the current state of affairs.

But sticking with the status quo can lead to boredom, which in turn can fuel complacency and stagnation. Borders, BlackBerry, Polaroid, and Myspace are but a few of the many companies that once had winning formulas but didn't update

their strategies until it was too late. Overly comfortable with the status quo, their leaders fell back on tradition and avoided the type of nonconformist behavior that could have spurred continued success.

We interpret information in a self-serving manner

A third reason for the prevalence of conformity is that we tend to prioritize information that supports our existing beliefs and to ignore information that challenges them, so we overlook things that could spur positive change. Complicating matters, we also tend to view unexpected or unpleasant information as a threat and to shun it—a phenomenon psychologists call *motivated skepticism*.

In fact, research suggests, the manner in which we weigh evidence resembles the manner in which we weigh ourselves on a bathroom scale. If the scale delivers bad news, we hop off and get back on—perhaps the scale misfired or we misread the display. If it delivers good news, we assume it's correct and cheerfully head for the shower.

Here's a more scientific example. Two psychologists, Peter Ditto and David Lopez, asked study participants to evaluate a student's intelligence by reviewing information about him one piece at a time—similar to the way college admissions officers evaluate applicants. The information was quite negative. Subjects could stop going through it as soon as they'd reached a firm conclusion. When they had been primed to like the

student (with a photo and some information provided before the evaluation), they turned over one card after another, searching for anything that would allow them to give a favorable rating. When they had been primed to dislike him, they turned over a few cards, shrugged, and called it a day.

By uncritically accepting information when it is consistent with what we believe and insisting on more when it isn’t, we subtly stack the deck against good decisions.

Promoting Constructive Nonconformity

Few leaders actively encourage deviant behavior in their employees; most go to great lengths to get rid of it. Yet nonconformity promotes innovation, improves performance, and can enhance a person’s standing more than conformity can. For example, research I conducted with Silvia Bellezza, of Columbia, and Anat Keinan, of Harvard, showed that observers judge a keynote speaker who wears red sneakers, a CEO who makes the rounds of Wall Street in a hoodie and jeans, and a presenter who creates her own PowerPoint template rather than using her company’s as having higher status than counterparts who conform to business norms.

My research also shows that going against the crowd gives us confidence in our actions, which makes us feel unique and engaged and translates to higher performance and greater creativity. In one field study, I asked a group of employees to

behave in nonconforming ways (speaking up if they disagreed with colleagues' decisions, expressing what they felt rather than what they thought they were expected to feel, and so on). I asked another group to behave in conforming ways, and a third group to do whatever its members usually did. After three weeks, those in the first group reported feeling more confident and engaged in their work than those in the other groups. They displayed more creativity in a task that was part of the study. And their supervisors gave them higher ratings on performance and innovativeness.

Six strategies can help leaders encourage constructive nonconformity in their organizations and themselves.

Step 1: Give Employees Opportunities to Be Themselves

Decades' worth of psychological research has shown that we feel accepted and believe that our views are more credible when our colleagues share them. But although conformity may make us feel good, it doesn't let us reap the benefits of authenticity. In one study Dan Cable, of London Business School, and Virginia Kay, then of the University of North Carolina at Chapel Hill, surveyed 154 recent MBA graduates who were four months into their jobs. Those who felt they could express their authentic selves at work were, on average, 16% more engaged and more committed to their organizations

than those who felt they had to hide their authentic selves. In another study, Cable and Kay surveyed 2,700 teachers who had been working for a year and reviewed the performance ratings given by their supervisors. Teachers who said they could express their authentic selves received higher ratings than teachers who did not feel they could do so.

Here are some ways to help workers be true to themselves:

Encourage employees to reflect on what makes them feel authentic. This can be done from the very start of the employment relationship—during orientation. In a field study I conducted with Brad Staats, of the University of North Carolina at Chapel Hill, and Dan Cable, employees in the business-process-outsourcing division of the Indian IT company Wipro went through a slightly modified onboarding process. We gave them a half hour to think about what was unique about them, what made them authentic, and how they could bring out their authentic selves at work. Later we compared them with employees who had gone through Wipro's usual onboarding program, which allowed no time for such reflection. The employees in the first group had found ways to tailor their jobs so that they could be their true selves—for example, they exercised judgment when answering calls instead of rigidly following the company script. They were more engaged in their work, performed better, and were more likely to be with the company seven months later.

The perils of conformity

Organizations put tremendous pressure on employees to conform. In a recent survey of 2,087 U.S. employees in a wide range of industries, nearly 49% agreed with the statement "I regularly feel pressure to conform in this organization."

This takes a heavy toll on individuals and enterprises alike. Employees who felt a need to conform reported a less positive work experience on several dimensions than did other employees, as shown by the average scores plotted below.

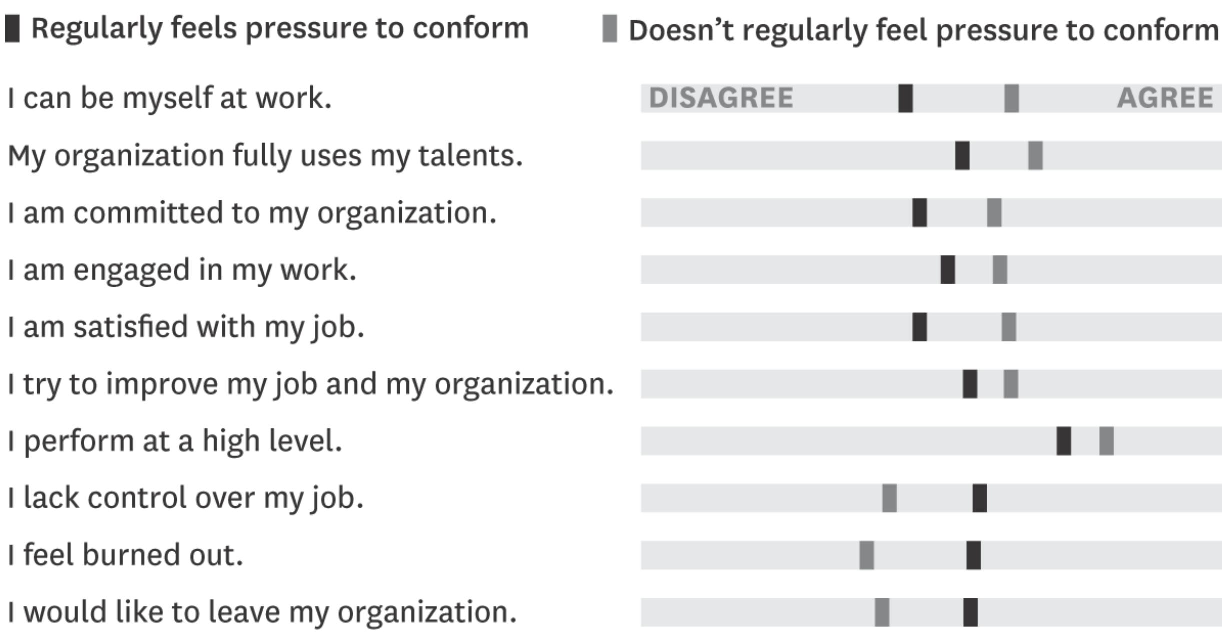

Leaders can also encourage this type of reflection once people are on the job. The start of a new year is a natural time for employees and their leaders to reflect on what makes them unique and authentic and how they can shape their jobs—even

in small ways—to avoid conformity. Reflection can also be encouraged at other career points, such as a performance review, a promotion, or a transition into a new role.

Tell employees what job needs to be done rather than how to do it. When Colleen Barrett was executive vice president of Southwest Airlines, from 1990 to 2001, she established the goal of allowing employees to be themselves. For example, flight attendants were encouraged to deliver the legally required safety announcement in their own style and with humor. “We have always thought that your avocation can be your vocation so that you don’t have to do any acting in your life when you leave home to go to work,” she has said. This philosophy helped make Southwest a top industry performer in terms of passenger volume, profitability, customer satisfaction, and turnover.

Let employees solve problems on their own. Leaders can encourage authenticity by allowing workers to decide how to handle certain situations. For instance, in the 1990s British Airways got rid of its thick customer-service handbook and gave employees the freedom (within reason) to figure out how to deal with customer problems as they arose (see “Competing on Customer Service: An Interview with British Airways’ Sir Colin Marshall,” HBR, November–December 1995).

Another company that subscribes to this philosophy is Pal's Sudden Service, a fast-food chain in the southern United States. By implementing lean principles, including the idea that workers are empowered to call out and fix problems, Pal's has achieved impressive numbers: one car served at the drive-through every 18 seconds, one mistake in every 3,600 orders (the industry average is one in 15), customer satisfaction scores of 98%, and health inspection scores above 97%. Turnover at the assistant manager level is under 2%, and in three decades Pal's has lost only seven general managers—two of them to retirement. Annual turnover on the front lines is about 34%—half the industry average. Pal's trains its employees extensively: New frontline workers receive 135 hours of instruction, on average (the industry average is about two hours). As a result, employees are confident that they can solve problems on their own and can stop processes if something does not seem right. (They also know they can ask for help.) When I was conducting interviews for a case on Pal's, a general manager gave me an example of how he encourages frontline workers to make decisions themselves: "A 16-year-old [employee] shows me a hot dog bun with flour on it and asks me if it's OK. My response: 'Your call. Would you sell it?'"

Let employees define their missions. Morning Star, a California-based tomato processing company, has employees write "personal commercial mission statements" that reflect who

they are and specify their goals for a given time period, ones that will contribute to the company's success. The statements are embedded in contracts known as "colleague letters of understanding," or CLOUs, which employees negotiate with coworkers, each spelling out how he or she will collaborate with others. The personal commercial mission of Morning Star's founder, Chris Rufer, is "to advance tomato technology to be the best in the world and operate these factories so they are pristine." That of one sales and marketing employee is "to indelibly mark 'Morning Star Tomato Products' on the tongue and brain of every commercial tomato product user." That of one employee in the shipping unit is "to reliably and efficiently provide our customers with marvelously attractive loads of desired product."

Step 2: Encourage Employees to Bring Out Their Signature Strengths

Michelangelo described sculpting as a process whereby the artist releases an ideal figure from the block of stone in which it slumbers. We all possess ideal forms, the signature strengths—being social connectors, for example, or being able to see the positive in any situation—that we use naturally in our lives. And we all have a drive to do what we do best and be recognized accordingly. A leader's task is to encourage

employees to sculpt their jobs to bring out their strengths—and to sculpt his or her own job, too. The actions below can help.

Give employees opportunities to identify their strengths. In a research project I conducted with Dan Cable, Brad Staats, and the University of Michigan's Julia Lee, leaders of national and local government agencies across the globe reflected each morning on their signature strengths and how to use them. They also read descriptions of times when they were at their best, written by people in their personal and professional networks. These leaders displayed more engagement and innovative behavior than members of a control group, and their teams performed better.

Tailor jobs to employees' strengths. Facebook is known for hiring smart people regardless of the positions currently open in the company, gathering information about their strengths, and designing their jobs accordingly. Another example is Osteria Francescana, a Michelin three-star restaurant in Modena, Italy, that won first place in the 2016 World's 50 Best Restaurant awards. Most restaurants, especially top-ranked ones, observe a strict hierarchy, with specific titles for each position. But at Osteria Francescana, jobs and their attendant responsibilities are tailored to individual workers.

Discovering employees' strengths takes time and effort. Massimo Bottura, the owner and head chef, rotates interns

through various positions for at least a few months so that he and his team can configure jobs to play to the newcomers' strengths. This ensures that employees land where they fit best.

If such a process is too ambitious for your organization, consider giving employees some freedom to choose responsibilities within their assigned roles.

Step 3: Question the Status Quo, and Encourage Employees to Do the Same

Although businesses can benefit from repeatable practices that ensure consistency, they can also stimulate employee engagement and innovation by questioning standard procedures—"the way we've always done it." Here are some proven tactics.

Ask "Why?" and "What if?" By regularly asking employees such questions, Max Zanardi, for several years the general manager of the Ritz-Carlton in Istanbul, creatively led them to redefine luxury by providing customers with authentic and unusual experiences. For example, employees had traditionally planted flowers each year on the terrace outside the hotel's restaurant. One day Zanardi asked, "Why do we always plant flowers? How about vegetables? What about herbs?" This resulted in a terrace garden featuring herbs and heirloom tomatoes used in the restaurant—things guests very much appreciated.

Leaders who question the status quo give employees reasons to stay engaged and often spark fresh ideas that can rejuvenate the business.

Stress that the company is not perfect. Ed Catmull, the cofounder and president of Pixar Animation Studios, worried that new hires would be too awed by Pixar's success to challenge existing practices (see "How Pixar Fosters Collective Creativity," HBR, September 2008). So during onboarding sessions, his speeches included examples of the company's mistakes. Emphasizing that we are all human and that the organization will never be perfect gives employees freedom to engage in constructive nonconformity.

Excel at the basics. Ensuring that employees have deep knowledge about the way things usually operate provides them with a foundation for constructively questioning the status quo. This philosophy underlies the many hours Pal's devotes to training: Company leaders want employees to be expert in all aspects of their work. Similarly, Bottura believes that to create innovative dishes, his chefs must be well versed in classic cooking techniques.

Step 4: Create Challenging Experiences

It's easy for workers to get bored and fall back on routine when their jobs involve little variety or challenge. And employees

who find their work boring lack the motivation to perform well and creatively, whereas work that is challenging enhances their engagement. Research led by David H. Zald, of Vanderbilt University, shows that novel behavior, such as trying something new or risky, triggers the release of dopamine, a chemical that helps keep us motivated and eager to innovate.

Leaders can draw on the following tactics when structuring employees' jobs:

Maximize variety. This makes it less likely that employees will go on autopilot and more likely that they will come up with innovative ways to improve what they're doing. It also boosts performance, as Brad Staats and I found in our analysis of two and a half years' worth of transaction data from a Japanese bank department responsible for processing home loan applications. The mortgage line involved 17 distinct tasks, including scanning applications, comparing scanned documents to originals, entering application data into the computer system, assessing whether information complied with underwriting standards, and conducting credit checks. Workers who were assigned diverse tasks from day to day were more productive than others (as measured by the time taken to complete each task); the variety kept them motivated. This allowed the bank to process applications more quickly, increasing its competitiveness.

Variety can be ensured in a number of ways. Pal's rotates employees through tasks (taking orders, grilling, working the register, and so on) in a different order each day. Some companies forgo defined career trajectories and instead move employees through various positions within departments or teams over the course of months or years.

In addition to improving engagement, job rotation broadens individuals' skill sets, creating a more flexible workforce. This makes it easier to find substitutes if someone falls ill or abruptly quits and to shift people from tasks where they are no longer needed (see "Why 'Good Jobs' Are Good for Retailers," HBR, January–February 2012).

Continually inject novelty into work. Novelty is a powerful force. When something new happens at work, we pay attention, engage, and tend to remember it. We are less likely to take our work for granted when it continues to generate strong feelings. Novelty in one's job is more satisfying than stability.

So, how can leaders inject it into work? Bottura throws last-minute menu changes at his team to keep excitement high. At Pal's, employees learn the order of their tasks for the day only when they get to work.

Leaders can also introduce novelty by making sure that projects include a few people who are somewhat out of their comfort zone, or by periodically giving teams new challenges (for instance, asking them to deliver a product faster than in

the past). They can assign employees to teams charged with designing a new work process or piloting a new service.

Identify opportunities for personal learning and growth. Giving people such experiences is an essential way to promote constructive nonconformity, research has shown. For instance, in a field study conducted at a global consulting firm, colleagues and I found that when onboarding didn't just focus on performance but also spotlighted opportunities for learning and growth, engagement and innovative behaviors were higher six months later. Companies often identify growth opportunities during performance reviews, of course, but there are many other ways to do so. Chefs at Osteria Francescana can accompany Bottura to cooking events that expose them to other countries, cuisines, traditions, arts, and culture—all potential sources of inspiration for new dishes. When I worked as a research consultant at Disney, in the summer of 2010, I learned that members of the Imagineering R&D group were encouraged to belong to professional societies, attend conferences, and publish in academic and professional journals. Companies can help pay for courses that may not strictly relate to employees' current jobs but would nonetheless expand their skill sets or fuel their curiosity.

Give employees responsibility and accountability. At Morning Star, if employees need new equipment to do their work—even

something that costs thousands of dollars—they may buy it. If they see a process that would benefit from different skills, they may hire someone. They must consult colleagues who would be affected (other people who would use the equipment, say), but they don't need approval from above. Because there are no job titles at Morning Star, how employees influence others—and thus get work done—is determined mainly by how their colleagues perceive the quality of their decisions.

Step 5: Foster Broader Perspectives

We often focus so narrowly on our own point of view that we have trouble understanding others' experiences and perspectives. And as we assume high-level positions, research shows, our egocentric focus becomes stronger. Here are some ways to combat it:

Create opportunities for employees to view problems from multiple angles. We all tend to be self-serving in terms of how we process information and generate (or fail to generate) alternatives to the status quo. Leaders can help employees overcome this tendency by encouraging them to view problems from different perspectives. At the electronics manufacturer Sharp, an oft-repeated maxim is "Be dragonflies, not flatfish." Dragonflies have compound eyes that can take in multiple perspectives at once; flatfish have both eyes on the same side of the head and can see in only one direction at a time.

Jon Olinto and Anthony Ackil, the founders of the fast-casual restaurant chain b.good, require all employees (including managers) and franchisees to be trained in every job—from prep to grill to register. (Unlike Pal's, however, b.good does not rotate people through jobs each day.) Being exposed to different perspectives increases engagement and innovative behaviors, research has found.

Use language that reduces self-serving bias. To prevent their traders from letting success go to their heads when the market is booming, some Wall Street firms regularly remind them, "Don't confuse brains with a bull market." At GE, terms such as "planting seeds" (to describe making investments that will produce fruitful results even after the managers behind them have moved on to other jobs) have entered the lexicon (see "How GE Teaches Teams to Lead Change," HBR, January 2009).

Hire people with diverse perspectives. Decades' worth of research has found that working among people from a variety of cultures and backgrounds helps us see problems in new ways and consider ideas that might otherwise go unnoticed, and it fosters the kind of creativity that champions change. At Osteria Francescana the two sous-chefs are Kondo "Taka" Takahiko, from Japan, and Davide di-Fabio, from Italy. They differ not only in country of origin but also in strengths and ways of thinking: Davide is comfortable with improvisation, for

example, while Taka is obsessed with precision. Diversity in ways of thinking is a quality sought by Rachael Chong, the founder and CEO of the startup Catchafire. When interviewing job candidates, she describes potential challenges and carefully listens to see whether people come up with many possible solutions or get stuck on a single one. To promote innovation and new approaches, Ed Catmull hires prominent outsiders, gives them important roles, and publicly acclaims their contributions. But many organizations do just the opposite: hire people whose thinking mirrors that of the current management team.

Step 6: Voice and Encourage Dissenting Views

We often seek out and fasten on information that confirms our beliefs. Yet data that is inconsistent with our views and may even generate negative feelings (such as a sense of failure) can provide opportunities to improve our organizations and ourselves. Leaders can use a number of tactics to push employees out of their comfort zones.

Look for disconfirming evidence. Leaders shouldn't ask, "Who agrees with this course of action?" or "What information supports this view?" Instead they should ask, "What information suggests this might not be the right path to take?" Mellody Hobson, the president of Ariel Investments and the chair of the board of directors of DreamWorks Animation,

regularly opens team meetings by reminding attendees that they don't need to be right; they need to bring up information that can help the team make the right decisions, which happens when members voice their concerns and disagree. At the Chicago Board of Trade, in-house investigators scrutinize trades that may violate exchange rules. To avoid bias in collecting information, they have been trained to ask open-ended interview questions, not ones that can be answered with a simple yes or no. Leaders can use a similar approach when discussing decisions. They should also take care not to depend on opinions but to assess whether the data supports or undermines the prevailing point of view.

Create dissent by default. Leaders can encourage debate during meetings by inviting individuals to take opposing points of view; they can also design processes to include dissent. When employees of Pal's suggest promising ideas for new menu items, the ideas are tested in three different stores: one whose owner-operator likes the idea ("the protagonist"), one whose owner-operator is skeptical ("the antagonist"), and one whose owner-operator has yet to form a strong opinion ("the neutral"). This ensures that dissenting views are aired and that they help inform the CEO's decisions about proposed items.

Identify courageous dissenters. Even if encouraged to push back, many timid or junior people won't. So make sure the team

includes people you know will voice their concerns, writes Diana McLain Smith in *The Elephant in the Room: How Relationships Make or Break the Success of Leaders and Organizations*. Once the more reluctant employees see that opposing views are welcome, they will start to feel comfortable dissenting as well.

Striking the Right Balance

By adopting the strategies above, leaders can fight their own and their employees' tendency to conform when that would hurt the company's interests. But to strike the optimal balance between conformity and nonconformity, they must think carefully about the boundaries within which employees will be free to deviate from the status quo. For instance, the way a manager leads her team can be up to her as long as her behavior is aligned with the company's purpose and values and she delivers on that purpose.

Assessment: Are You a "Constructive Nonconformist"?

Find out how much of a rebel worker you are.

For decades, prevailing management wisdom has encouraged leaders to focus on designing efficient processes and getting employees to follow them. But conformity can hurt businesses. Innovation and high performance often result

from behaviors that defy organizational norms—established ways of thinking and of doing things. How much does your company pressure you to conform? And are you succumbing to the pressure and hurting your chances of success? Take the following assessment (adapted from my ongoing research) to discover whether you're engaging in what I call constructive nonconformity: deviant behavior that benefits the organization.

When answering these questions, focus on the past month.	***Never***	***Almost never***	***Sometimes***	***Fairly often***	***Very often***	***Always***
1. In the past month, how often have you refrained from opposing your team members just to avoid rocking the boat?	0	1	2	3	4	5
2. How often have you publicly supported ideas you privately disagreed with?	0	1	2	3	4	5
3. How often have you followed established rules or procedures, even though you suspected there was a better way to do things?	0	1	2	3	4	5
4. How often have you raised questions about the effectiveness of current processes or systems?	5	4	3	2	1	0
5. How often have you seen senior leaders challenge the status quo or ask employees to think outside the box?	5	4	3	2	1	0
6. How often have you felt pressured to conform to the cultural norms of your organization (how to dress, how to interact with others, how to do your work, and so on)?	0	1	2	3	4	5

7. How often have you felt free to be yourself—to behave and express yourself in an authentic way?	5	4	3	2	1	0
8. How often have you been encouraged to solve problems on your own, without involving a supervisor?	5	4	3	2	1	0
9. How often has your job played to your strengths?	5	4	3	2	1	0
10. How often have you been challenged—urged to develop a new skill or to take on a task that pushed you out of your comfort zone?	5	4	3	2	1	0
11. How often have you sought information that was inconsistent with your views and might even prove you wrong?	5	4	3	2	1	0
12. How often have you and your team been encouraged to debate ideas or consider multiple perspectives before reaching a decision?	5	4	3	2	1	0

Score: 0–24 You're lucky: Your low score indicates that you are probably very engaged in your work, are performing at a high level, and are innovating frequently. Just make sure that you don't become complacent—the pressure to conform affects everyone. Keep being the rebel that you are!

Score: 25–30 Your score is average—and in this case, average is good. Scores in this range indicate that your ability to express yourself at work is at a healthy level, allowing you to be productive and innovative. To stay in this sweet spot, watch out for situations in which you feel pressured to conform.

Score: 31–39 Your higher-than-average score indicates a level of pressure that may be detrimental to your performance and your ability to innovate. You may also be disengaged. Try shaping your job in ways that allow you to be yourself and that bring out your talents and skills. Even small changes can let your authentic self shine through.

Score: 40–60 Your high score indicates an unproductive level of conformity. You're probably disengaged, and you're almost certainly having a hard time being your true self at work. It's critical that you find ways (big and small) to lower the pressure to conform, and that starts with allowing your authentic self to shine through. Act more like a rebel, and you and your organization will benefit.

Morning Star's colleague letters of understanding provide such boundaries. They clearly state employees' goals and their responsibility to deliver on the organization's purpose but leave it up to individual workers to decide how to achieve those goals. Colleagues with whom an employee has negotiated a CLOU will let him know if his actions cross a line.

Brazil's Semco Group, a 3,000-employee conglomerate, similarly relies on peer pressure and other mechanisms to give employees considerable freedom while making sure they don't go overboard. The company has no job titles, dress code, or organizational charts. If you need a workspace, you reserve it in one of a few satellite offices scattered around São Paulo. Employees, including factory workers, set their own schedules and production quotas. They even choose the amount and

form of their compensation. What prevents employees from taking advantage of this freedom? First, the company believes in transparency: All its financial information is public, so everyone knows what everyone else makes. People who pay themselves too much have to work with resentful colleagues. Second, employee compensation is tied directly to company profits, creating enormous peer pressure to keep budgets in line.

Ritz-Carlton, too, excels in balancing conformity and nonconformity. It depends on 3,000 standards developed over the years to ensure a consistent customer experience at all its hotels. These range from how to slice a lime to which toiletries to stock in the bathrooms. But employees have considerable freedom within those standards and can question them if they see ways to provide an even better customer experience. For instance, for many years the company has allowed staff members to spend up to $2,000 to address any customer complaint in the way they deem best. (Yes, that is $2,000 per employee per guest.) The hotel believes that business is most successful when employees have well-defined standards, understand the reasoning behind them, and are given autonomy in carrying them out.

Organizations, like individuals, can easily become complacent, especially when business is going well. Complacency often sets in because of too much conformity—stemming from peer pressure, acceptance of the status quo,

and the interpretation of information in self-serving ways. The result is a workforce of people who feel they can't be themselves on the job, are bored, and don't consider others' points of view.

Constructive nonconformity can help companies avoid these problems. If leaders were to put just half the time they spend ensuring conformity into designing and installing mechanisms to encourage constructive deviance, employee engagement, productivity, and innovation would soar.

Further Reading

IN THE COURSE OF DEVELOPING this Big Idea on Rebel Talent, HBR asked Francesca Gino to provide a portfolio of content that could further inspire, advise, and help develop your understanding of the topic. Gino's curated list of materials on rebel talent runs the gamut from classic HBR articles to novels and more.

HBR Articles

While studying leaders and organizations that attract, develop, and manage talent so as to spark engagement and creativity, I found many insights in the pages of HBR.

- **"How Pixar Fosters Collective Creativity,"** Ed Catmull, September 2008

- **"Are You a High Potential?,"** Douglas A. Ready, Jay A. Conger, and Linda A. Hill, June 2010
- **"How to Hang On to Your High Potentials,"** Claudio Fernández-Aráoz, Boris Groysberg, and Nitin Nohria, October 2011
- **"How GE Teaches Teams to Lead Change,"** Steven Prokesch, January 2009
- **"Managing Without Managers,"** Ricardo Semler, September–October 1989
- **"Why My Former Employees Still Work for Me,"** Ricardo Semler, January–February 1994

Books

I've found inspiration in books from as far back as the 1950s that document how and why companies create pressure to conform and what can be done to combat it.

- ***The Organization Man,*** William H. Whyte, 1956
- ***Reinventing Organizations: A Guide to Creating Organizations Inspired by the Next Stage of Human Consciousness,*** Frederic Laloux, 2014
- ***The Art of Being Unmistakable: A Collection of Essays About Making a Dent in the Universe,*** Srinivas Rao, 2013
- ***Bartleby, the Scrivener,*** Herman Melville, 1853

- ***Collective Genius: The Art and Practice of Leading Innovation,*** Linda A. Hill, Greg Brandeau, Emily Truelove, and Kent Lineback, 2014

Case Studies

The best way to learn how to foster constructive nonconformity is to dig into how actual companies did so.

- ***"Sun Hydraulics: Leading in Tough Times (A),"*** Linda A. Hill and Jennifer M. Suesse, 2003
- ***"Pal's Sudden Service—Scaling an Organizational Model to Drive Growth,"*** Gary P. Pisano, Francesca Gino, and Bradley R. Staats, 2016
- ***"The Morning Star Company: Self-Management at Work,"*** Francesca Gino and Bradley R. Staats, 2013

Other Articles

- ***"Monkeys Are Adept at Picking Up Social Cues, Research Shows,"*** Pam Belluck, *New York Times*, 2013
- ***"For Some Flight Attendants, Shtick Comes With the Safety Spiel,"*** Zach Schonbrun, *New York Times*, 2016
- ***"I'm Quite Eccentric Within Accepted Societal Norms,"*** Martin Grossman, *The Onion*, 2007

Originally published in October–November 2016. Reprint H035GG

Why Diversity Programs Fail

by Frank Dobbin and Alexandra Kalev

BUSINESSES STARTED CARING A LOT more about diversity after a series of high-profile lawsuits rocked the financial industry. In the late 1990s and early 2000s, Morgan Stanley shelled out $54 million—and Smith Barney and Merrill Lynch more than $100 million each—to settle sex discrimination claims. In 2007, Morgan was back at the table, facing a new class action, which cost the company $46 million. In 2013, Bank of America Merrill Lynch settled a race discrimination suit for $160 million. Cases like these brought Merrill's total 15-year payout to nearly *half a billion* dollars.

It's no wonder that Wall Street firms now require new hires to sign arbitration contracts agreeing not to join class actions. They have also expanded training and other diversity programs. But on balance, equality isn't improving in financial services or elsewhere. Although the proportion of managers at U.S. commercial banks who were Hispanic rose from 4.7% in 2003 to 5.7% in 2014, white women's representation dropped

from 39% to 35%, and black men's from 2.5% to 2.3%. The numbers were even worse in investment banks (though that industry is shrinking, which complicates the analysis). Among all U.S. companies with 100 or more employees, the proportion of black men in management increased just slightly—from 3% to 3.3%—from 1985 to 2014. White women saw bigger gains from 1985 to 2000—rising from 22% to 29% of managers—but their numbers haven't budged since then. Even in Silicon Valley, where many leaders tout the need to increase diversity for both business and social justice reasons, bread-and-butter tech jobs remain dominated by white men.

It shouldn't be surprising that most diversity programs aren't increasing diversity. Despite a few new bells and whistles, courtesy of big data, companies are basically doubling down on the same approaches they've used since the 1960s—which often make things worse, not better. Firms have long relied on diversity training to reduce bias on the job, hiring tests and performance ratings to limit it in recruitment and promotions, and grievance systems to give employees a way to challenge managers. Those tools are designed to preempt lawsuits by policing managers' thoughts and actions. Yet laboratory studies show that this kind of force-feeding can activate bias rather than stamp it out. As social scientists have found, people often rebel against rules to assert their autonomy. Try to coerce me to do X, Y, or Z, and I'll do the opposite just to prove that I'm my own person.

In analyzing three decades' worth of data from more than 800 U.S. firms and interviewing hundreds of line managers and executives at length, we've seen that companies get better results when they ease up on the control tactics. It's more effective to engage managers in solving the problem, increase their on-the-job contact with female and minority workers, and promote social accountability—the desire to look fair-minded. That's why interventions such as targeted college recruitment, mentoring programs, self-managed teams, and task forces have boosted diversity in businesses. Some of the most effective solutions aren't even designed with diversity in mind.

Here, we dig into the data, the interviews, and company examples to shed light on what doesn't work and what does.

Why You Can't Just Outlaw Bias

Executives favor a classic command-and-control approach to diversity because it boils expected behaviors down to dos and don'ts that are easy to understand and defend. Yet this approach also flies in the face of nearly everything we know about how to motivate people to make changes. Decades of social science research point to a simple truth: You won't get managers on board by blaming and shaming them with rules and reeducation. Let's look at how the most common top-down efforts typically go wrong.

Idea in Brief

The Problem

To reduce bias and increase diversity, organizations are relying on the same programs they've been using since the 1960s. Some of these efforts make matters worse, not better.

The Reason

Most diversity programs focus on controlling managers' behavior, and as studies show, that approach tends to activate bias rather than quash it. People rebel against rules that threaten their autonomy.

The Solution

Instead of trying to police managers' decisions, the most effective programs engage people in working for diversity, increase their contact with women and minorities, and tap into their desire to look good to others.

Diversity training

Do people who undergo training usually shed their biases? Researchers have been examining that question since before World War II, in nearly a thousand studies. It turns out that while people are easily taught to respond correctly to a questionnaire about bias, they soon forget the right answers. The positive effects of diversity training rarely last beyond a day or two, and a number of studies suggest that it can activate

bias or spark a backlash. Nonetheless, nearly half of midsize companies use it, as do nearly all the *Fortune* 500.

Many firms see adverse effects. One reason is that three-quarters use negative messages in their training. By headlining the legal case for diversity and trotting out stories of huge settlements, they issue an implied threat: “Discriminate, and the company will pay the price.” We understand the temptation—that’s how we got your attention in the first paragraph—but threats, or “negative incentives,” don’t win converts.

Another reason is that about three-quarters of firms with training still follow the dated advice of the late diversity guru R. Roosevelt Thomas Jr. “If diversity management is strategic to the organization,” he used to say, diversity training must be mandatory, and management has to make it clear that “if you can’t deal with that, then we have to ask you to leave.” But five years after instituting required training for managers, companies saw no improvement in the proportion of white women, black men, and Hispanics in management, and the share of black women actually decreased by 9%, on average, while the ranks of Asian-American men and women shrank by 4% to 5%. Trainers tell us that people often respond to compulsory courses with anger and resistance—and many participants actually report more animosity toward other groups afterward.

But voluntary training evokes the opposite response (“I chose to show up, so I must be pro-diversity”), leading to better results: increases of 9% to 13% in black men, Hispanic men, and Asian-American men and women in management five years out (with no decline in white or black women). Research from the University of Toronto reinforces our findings: In one study white subjects read a brochure critiquing prejudice toward blacks. When people felt pressure to agree with it, the reading strengthened their bias against blacks. When they felt the choice was theirs, the reading reduced bias.

Companies too often signal that training is remedial. The diversity manager at a national beverage company told us that the top brass uses it to deal with problem groups. “If there are a number of complaints … or, God forbid, some type of harassment case … leaders say, ‘Everyone in the business unit will go through it again.’” Most companies with training have special programs for managers. To be sure, they’re a high-risk group because they make the hiring, promotion, and pay decisions. But singling them out implies that they’re the worst culprits. Managers tend to resent that implication and resist the message.

Hiring tests

Some 40% of companies now try to fight bias with mandatory hiring tests assessing the skills of candidates for frontline jobs. But managers don’t like being told that they can’t hire

whomever they please, and our research suggests that they often use the tests selectively. Back in the 1950s, following the postwar migration of blacks northward, Swift & Company, Chicago meatpackers, instituted tests for supervisor and quality-checking jobs. One study found managers telling blacks that they had failed the test and then promoting whites who hadn't been tested. A black machine operator reported: "I had four years at Englewood High School. I took an exam for a checker's job. The foreman told me I failed" and gave the job to a white man who "didn't take the exam."

This kind of thing still happens. When we interviewed the new HR director at a West Coast food company, he said he found that white managers were making only strangers—most of them minorities—take supervisor tests and hiring white friends without testing them. "If you are going to test one person for this particular job title," he told us, "you need to test everybody."

But even managers who test everyone applying for a position may ignore the results. Investment banks and consulting firms build tests into their job interviews, asking people to solve math and scenario-based problems on the spot. While studying this practice, Kellogg professor Lauren Rivera played a fly on the wall during hiring meetings at one firm. She found that the team paid little attention when white men blew the math test but close attention when women and blacks did. Because

decision makers (deliberately or not) cherry-picked results, the testing amplified bias rather than quashed it.

Companies that institute written job tests for managers—about 10% have them today—see decreases of 4% to 10% in the share of managerial jobs held by white women, African-American men and women, Hispanic men and women, and Asian-American women over the next five years. There are significant declines among white and Asian-American women—groups with high levels of education, which typically score well on standard managerial tests. So group differences in test-taking skills don't explain the pattern.

Performance ratings

More than 90% of midsize and large companies use annual performance ratings to ensure that managers make fair pay and promotion decisions. Identifying and rewarding the best workers isn't the only goal—the ratings also provide a litigation shield. Companies sued for discrimination often claim that their performance rating systems prevent biased treatment.

But studies show that raters tend to lowball women and minorities in performance reviews. And some managers give everyone high marks to avoid hassles with employees or to keep their options open when handing out promotions. However managers work around performance systems, the bottom line is that ratings don't boost diversity. When companies introduce them, there's no effect on minority

managers over the next five years, and the share of white women in management drops by 4%, on average.

Grievance procedures

This last tactic is meant to identify and rehabilitate biased managers. About half of midsize and large firms have systems through which employees can challenge pay, promotion, and termination decisions. But many managers—rather than change their own behavior or address discrimination by others—try to get even with or belittle employees who complain. Among the nearly 90,000 discrimination complaints made to the Equal Employment Opportunity Commission in 2015, 45% included a charge of retaliation—which suggests that the original report was met with ridicule, demotion, or worse.

Once people see that a grievance system isn't warding off bad behavior in their organization, they may become less likely to speak up. Indeed, employee surveys show that most people don't report discrimination. This leads to another unintended consequence: Managers who receive few complaints conclude that their firms don't have a problem. We see this a lot in our interviews. When we talked with the vice president of HR at an electronics firm, she mentioned the widely publicized "difficulties other corporations are having" and added, "We have not had any of those problems … we have gone almost four years without any kind of discrimination complaint!" What's more, lab studies show that protective measures like

grievance systems lead people to drop their guard and let bias affect their decisions, because they think company policies will guarantee fairness.

Things don't get better when firms put in formal grievance systems; they get worse. Our quantitative analyses show that the managerial ranks of white women and all minority groups except Hispanic men decline—by 3% to 11%—in the five years after companies adopt them.

Still, most employers feel they need some sort of system to intercept complaints, if only because judges like them. One strategy that is gaining ground is the "flexible" complaint system, which offers not only a formal hearing process but also informal mediation. Since an informal resolution doesn't involve hauling the manager before a disciplinary body, it may reduce retaliation. As we'll show, making managers feel accountable without subjecting them to public rebuke tends to help.

Tools for Getting Managers on Board

If these popular solutions backfire, then what can employers do instead to promote diversity?

A number of companies have gotten consistently positive results with tactics that don't focus on control. They apply three basic principles: engage managers in solving the problem,

expose them to people from different groups, and encourage social accountability for change.

Engagement

When someone's beliefs and behavior are out of sync, that person experiences what psychologists call "cognitive dissonance." Experiments show that people have a strong tendency to "correct" dissonance by changing either the beliefs or the behavior. So, if you prompt them to act in ways that support a particular view, their opinions shift toward that view. Ask them to write an essay defending the death penalty, and even the penalty's staunch opponents will come to see some merits. When managers actively help boost diversity in their companies, something similar happens: They begin to think of themselves as diversity champions.

Take *college recruitment programs* targeting women and minorities. Our interviews suggest that managers willingly participate when invited. That's partly because the message is positive: "Help us find a greater variety of promising employees!" And involvement is voluntary: Executives sometimes single out managers they think would be good recruiters, but they don't drag anyone along at gunpoint.

Managers who make college visits say they take their charge seriously. They are determined to come back with strong candidates from underrepresented groups—female engineers, for instance, or African-American management trainees.

Cognitive dissonance soon kicks in—and managers who were wishy-washy about diversity become converts.

The effects are striking. Five years after a company implements a college recruitment program targeting female employees, the share of white women, black women, Hispanic women, and Asian-American women in its management rises by about 10%, on average. A program focused on minority recruitment increases the proportion of black male managers by 8% and black female managers by 9%.

Mentoring is another way to engage managers and chip away at their biases. In teaching their protégés the ropes and sponsoring them for key training and assignments, mentors help give their charges the breaks they need to develop and advance. The mentors then come to believe that their protégés merit these opportunities—whether they're white men, women, or minorities. That is cognitive dissonance—"Anyone I sponsor must be deserving"—at work again.

While white men tend to find mentors on their own, women and minorities more often need help from formal programs. One reason, as Georgetown's business school dean David Thomas discovered in his research on mentoring, is that white male executives don't feel comfortable reaching out informally to young women and minority men. Yet they are eager to mentor assigned protégés, and women and minorities are often first to sign up for mentors.

Mentoring programs make companies' managerial echelons significantly more diverse: On average they boost the representation of black, Hispanic, and Asian-American women, and Hispanic and Asian-American men, by 9% to 24%. In industries where plenty of college-educated nonmanagers are eligible to move up, like chemicals and electronics, mentoring programs also increase the ranks of white women and black men by 10% or more.

Only about 15% of firms have special college recruitment programs for women and minorities, and only 10% have mentoring programs. Once organizations try them out, though, the upside becomes clear. Consider how these programs helped Coca-Cola in the wake of a race discrimination suit settled in 2000 for a record $193 million. With guidance from a court-appointed external task force, executives in the North America group got involved in recruitment and mentoring initiatives for professionals and middle managers, working specifically toward measurable goals for minorities. Even top leaders helped to recruit and mentor, and talent-sourcing partners were required to broaden their recruitment efforts. After five years, according to former CEO and chairman Neville Isdell, 80% of all mentees had climbed at least one rung in management. Both individual and group mentoring were open to all races but attracted large numbers of African-Americans (who accounted for 36% of protégés). These changes brought important gains. From 2000 to 2006, African-Americans'

representation among salaried employees grew from 19.7% to 23%, and Hispanics' from 5.5% to 6.4%. And while African-Americans and Hispanics respectively made up 12% and 4.9% of professionals and middle managers in 2002, just four years later those figures had risen to 15.5% and 5.9%.

This began a virtuous cycle. Today, Coke looks like a different company. This February, *Atlanta Tribune* magazine profiled 17 African-American women in VP roles and above at Coke, including CFO Kathy Waller.

Contact

Evidence that contact between groups can lessen bias first came to light in an unplanned experiment on the European front during World War II. The U.S. army was still segregated, and only whites served in combat roles. High casualties left General Dwight Eisenhower understaffed, and he asked for black volunteers for combat duty. When Harvard sociologist Samuel Stouffer, on leave at the War Department, surveyed troops on their racial attitudes, he found that whites whose companies had been joined by black platoons showed dramatically lower racial animus and greater willingness to work alongside blacks than those whose companies remained segregated. Stouffer concluded that whites fighting alongside blacks came to see them as soldiers like themselves first and foremost. The key, for Stouffer, was that whites and blacks had to be working toward a common goal *as equals*—hundreds of

years of close contact during and after slavery hadn't dampened bias.

Business practices that generate this kind of contact across groups yield similar results. Take *self-managed teams,* which allow people in different roles and functions to work together on projects as equals. Such teams increase contact among diverse types of people, because specialties within firms are still largely divided along racial, ethnic, and gender lines. For example, women are more likely than men to work in sales, whereas white men are more likely to be in tech jobs and management, and black and Hispanic men are more likely to be in production.

As in Stouffer's combat study, working side-by-side breaks down stereotypes, which leads to more equitable hiring and promotion. At firms that create self-managed work teams, the share of white women, black men and women, and Asian-American women in management rises by 3% to 6% over five years.

Rotating management trainees through departments is another way to increase contact. Typically, this kind of *cross-training* allows people to try their hand at various jobs and deepen their understanding of the whole organization. But it also has a positive impact on diversity, because it exposes both department heads and trainees to a wider variety of people. The result, we've seen, is a bump of 3% to 7% in white women,

black men and women, and Asian-American men and women in management.

About a third of U.S. firms have self-managed teams for core operations, and nearly four-fifths use cross-training, so these tools are already available in many organizations. Though college recruitment and mentoring have a bigger impact on diversity—perhaps because they activate engagement in the diversity mission *and* create intergroup contact—every bit helps. Self-managed teams and cross-training have had more positive effects than mandatory diversity training, performance evaluations, job testing, or grievance procedures, which are supposed to promote diversity.

Social accountability

The third tactic, encouraging social accountability, plays on our need to look good in the eyes of those around us. It is nicely illustrated by an experiment conducted in Israel. Teachers in training graded identical compositions attributed to Jewish students with Ashkenazic names (European heritage) or with Sephardic names (African or Asian heritage). Sephardic students typically come from poorer families and do worse in school. On average, the teacher trainees gave the Ashkenazic essays Bs and the Sephardic essays Ds. The difference evaporated, however, when trainees were told that they would discuss their grades with peers. The idea that they might have

to explain their decisions led them to judge the work by its quality.

In the workplace you'll see a similar effect. Consider this field study conducted by Emilio Castilla of MIT's Sloan School of Management: A firm found it consistently gave African-Americans smaller raises than whites, even when they had identical job titles and performance ratings. So Castilla suggested transparency to activate social accountability. The firm posted each unit's average performance rating and pay raise by race and gender. Once managers realized that employees, peers, and superiors would know which parts of the company favored whites, the gap in raises all but disappeared.

Corporate *diversity task forces* help promote social accountability. CEOs usually assemble these teams, inviting department heads to volunteer and including members of underrepresented groups. Every quarter or two, task forces look at diversity numbers for the whole company, for business units, and for departments to figure out what needs attention.

After investigating where the problems are—recruitment, career bottlenecks, and so on—task force members come up with solutions, which they then take back to their departments. They notice if their colleagues aren't volunteering to mentor or showing up at recruitment events. Accountability theory suggests that having a task force member in a department will cause managers in it to ask

themselves, “Will this look right?” when making hiring and promotion decisions.

Which Diversity Efforts Actually Succeed?

IN 829 MIDSIZE AND LARGE U.S. FIRMS, we analyzed how various diversity initiatives affected the proportion of women and minorities in management. Here you can see which ones helped different groups gain ground—and which set them back, despite good intentions. (No bar means we can’t say with statistical certainty if the program had any effect.)

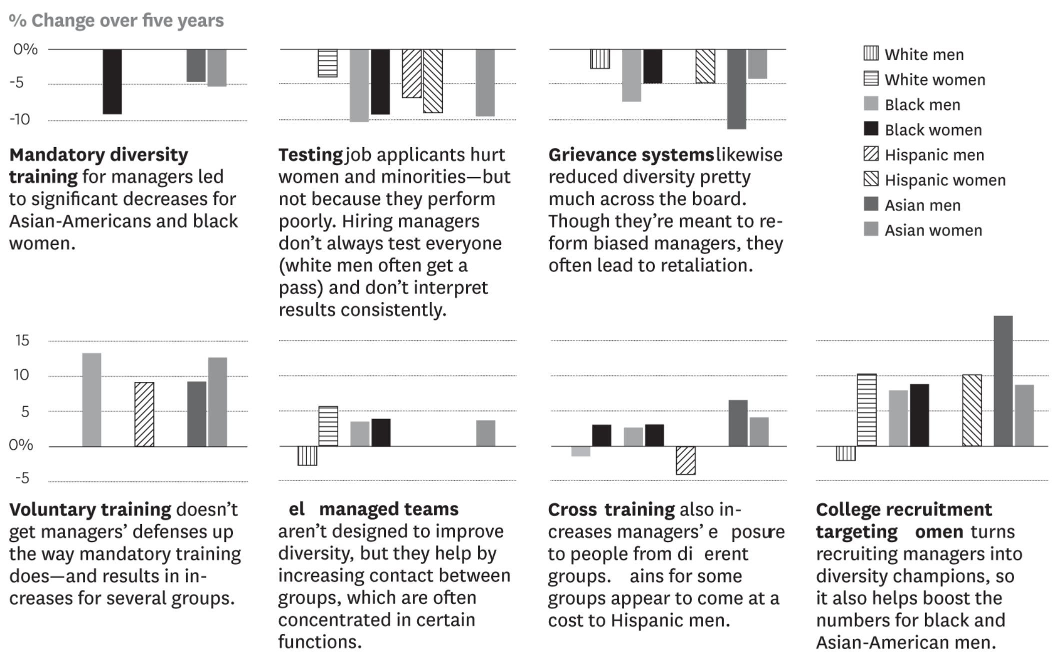

Poor returns on the usual programs

The three most popular interventions made firms less diverse, not more, because managers resisted strong-arming.

Programs that get results

Companies do a better job of increasing diversity when they forgo the control tactics and frame their efforts more positively. The most effective programs spark engagement, increase contact among different groups, or draw on people's strong desire to look good to others.

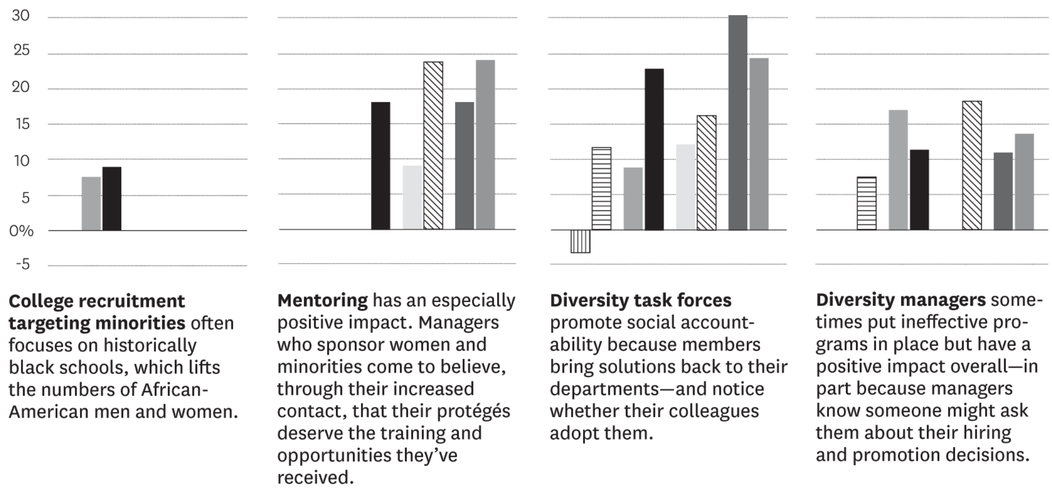

Note: In our analysis, we've isolated the effects of diversity programs from everything else going on in the companies and in the economy.

Deloitte has seen how powerful social accountability can be. In 1992, Mike Cook, who was then the CEO, decided to try to stanch the hemorrhaging of female associates. Half the company's hires were women, but nearly all of them left before they were anywhere near making partner. As Douglas McCracken, CEO of Deloitte's consulting unit at the time, later recounted in HBR, Cook assembled a high-profile task force that "didn't immediately launch a slew of new organizational policies aimed at outlawing bad behavior" but, rather, relied on transparency to get results.

The task force got each office to monitor the career progress of its women and set its own goals to address local problems. When it became clear that the CEO and other managing partners were closely watching, McCracken wrote, “women started getting their share of premier client assignments and informal mentoring.” And unit heads all over the country began getting questions from partners and associates about why things weren’t changing faster. An external advisory council issued annual progress reports, and individual managers chose change metrics to add to their own performance ratings. In eight years turnover among women dropped to the same level as turnover among men, and the proportion of female partners increased from 5% to 14%—the highest percentage among the big accounting firms. By 2015, 21% of Deloitte’s global partners were women, and in March of that year, Deloitte LLP appointed Cathy Engelbert as its CEO—making her the first woman to head a major accountancy.

Task forces are the trifecta of diversity programs. In addition to promoting accountability, they engage members who might have previously been cool to diversity projects and increase contact among the women, minorities, and white men who participate. They pay off, too: On average, companies that put in diversity task forces see 9% to 30% increases in the representation of white women and of each minority group in management over the next five years.

Diversity managers, too, boost inclusion by creating social accountability. To see why, let's go back to the finding of the teacher-in-training experiment, which is supported by many studies: When people know they *might* have to explain their decisions, they are less likely to act on bias. So simply having a diversity manager who could ask them questions prompts managers to step back and consider everyone who is qualified instead of hiring or promoting the first people who come to mind. Companies that appoint diversity managers see 7% to 18% increases in all underrepresented groups—except Hispanic men—in management in the following five years. Those are the gains after accounting for both effective and ineffective programs they put in place.

Only 20% of medium and large employers have task forces, and just 10% have diversity managers, despite the benefits of both. Diversity managers cost money, but task forces use existing workers, so they're a lot cheaper than some of the things that fail, such as mandatory training.

Leading companies like Bank of America Merrill Lynch, Facebook, and Google have placed big bets on accountability in the past couple of years. Expanding on Deloitte's early example, they're now posting complete diversity numbers for all to see. We should know in a few years if that moves the needle for them.

Strategies for controlling bias—which drive most diversity efforts—have failed spectacularly since they were introduced to promote equal opportunity. Black men have barely gained ground in corporate management since 1985. White women haven't progressed since 2000. It isn't that there aren't enough educated women and minorities out there—both groups have made huge educational gains over the past two generations. The problem is that we can't motivate people by forcing them to get with the program and punishing them if they don't.

The numbers sum it up. Your organization will become less diverse, not more, if you require managers to go to diversity training, try to regulate their hiring and promotion decisions, and put in a legalistic grievance system.

The very good news is that we know what does work—we just need to do more of it.

Originally published in July–August 2016. Reprint R1607C

What So Many People Don't Get About the U.S. Working Class

by Joan C. Williams

MY FATHER-IN-LAW GREW UP eating blood soup. He hated it, whether because of the taste or the humiliation, I never knew. His alcoholic father regularly drank up the family wage, and the family was often short on food money. They were evicted from apartment after apartment.

He dropped out of school in eighth grade to help support the family. Eventually he got a good, steady job he truly hated, as an inspector in a factory that made those machines that measure humidity levels in museums. He tried to open several businesses on the side but none worked, so he kept that job for 38 years. He rose from poverty to a middle-class life: the car, the house, two kids in Catholic school, the wife who worked only part-time. He worked incessantly. He had two jobs in addition to his full-time position, one doing yard work for a local magnate and another hauling trash to the dump.

Throughout the 1950s and 1960s, he read the *Wall Street Journal* and voted Republican. He was a man before his time: a blue-collar white man who thought the union was a bunch of jokers who took your money and never gave you anything in return. Starting in 1970, many blue-collar whites followed his example. This week, their candidate won the presidency.

For months, the only thing that's surprised me about Donald Trump is my friends' astonishment at his success. What's driving it is the class culture gap.

One little-known element of that gap is that the white working class (WWC) resents professionals but admires the rich. Class migrants (white-collar professionals born to blue-collar families) report that "professional people were generally suspect" and that managers are college kids "who don't know shit about how to do anything but are full of ideas about how I have to do my job," said Alfred Lubrano in *Limbo.* Barbara Ehrenreich recalled in 1990 that her blue-collar dad "could not say the word *doctor* without the virtual prefix *quack.* Lawyers were *shysters* … and professors were without exception *phonies.*" Annette Lareau found tremendous resentment against teachers, who were perceived as condescending and unhelpful.

Michèle Lamont, in *The Dignity of Working Men,* also found resentment of professionals—but not of the rich. "[I] can't knock anyone for succeeding," a laborer told her. "There's a lot of people out there who are wealthy and I'm sure they worked darned hard for every cent they have," chimed in a receiving

clerk. Why the difference? For one thing, most blue-collar workers have little direct contact with the rich outside of *Lifestyles of the Rich and Famous.* But professionals order them around every day. The dream is not to become upper-middle-class, with its different food, family, and friendship patterns; the dream is to live in your own class milieu, where you feel comfortable—just with more money. "The main thing is to be independent and give your own orders and not have to take them from anybody else," a machine operator told Lamont. Owning one's own business—that's the goal. That's another part of Trump's appeal.

Hillary Clinton, by contrast, epitomizes the dorky arrogance and smugness of the professional elite. The dorkiness: the pantsuits. The arrogance: the email server. The smugness: the basket of deplorables. Worse, her mere presence rubs it in that *even women* from her class can treat working-class men with disrespect. Look at how she condescends to Trump as unfit to hold the office of the presidency and dismisses his supporters as racist, sexist, homophobic, or xenophobic.

Trump's blunt talk taps into another blue-collar value: straight talk. "Directness is a working-class norm," notes Lubrano. As one blue-collar guy told him, "If you have a problem with me, come talk to me. If you have a way you want something done, come talk to me. I don't like people who play these two-faced games." Straight talk is seen as requiring manly courage, not being "a total wuss and a wimp," an electronics

technician told Lamont. Of course Trump appeals. Clinton's clunky admission that she talks one way in public and another in private? Further proof she's a two-faced phony.

Manly dignity is a big deal for working-class men, and they're not feeling that they have it. Trump promises a world free of political correctness and a return to an earlier era, when men were men and women knew their place. It's comfort food for high-school-educated guys who could have been my father-in-law if they'd been born 30 years earlier. Today they feel like losers—or did until they met Trump.

Manly dignity is a big deal for most men. So is breadwinner status: Many still measure masculinity by the size of a paycheck. White working-class men's wages hit the skids in the 1970s and took another body blow during the Great Recession. Look, I wish manliness worked differently. But most men, like most women, seek to fulfill the ideals they've grown up with. For many blue-collar men, all they're asking for is basic human dignity (male varietal). Trump promises to deliver it.

The Democrats' solution? Last week the *New York Times* published an article advising men with high-school educations to take pink-collar jobs. Talk about insensitivity. Elite men, you will notice, are not flooding into traditionally feminine work. To recommend that for WWC men just fuels class anger.

Isn't what happened to Clinton unfair? Of course it is. It is unfair that she wasn't a plausible candidate until she was so overqualified she was suddenly unqualified due to past

mistakes. It is unfair that Clinton is called a “nasty woman” while Trump is seen as a real man. It’s unfair that Clinton only did so well in the first debate because she wrapped her candidacy in a shimmy of femininity. When she returned to attack mode, it was the right thing for a presidential candidate to do but the wrong thing for a woman to do. The election shows that sexism retains a deeper hold than most imagined. But women don’t stand together: WWC women voted for Trump over Clinton by a whopping 28-point margin—62% to 34%. If they’d split 50-50, she would have won.

Class trumps gender, and it’s driving American politics. Policy makers of both parties—but particularly Democrats if they are to regain their majorities—need to remember five major points.

Understand That Working Class Means Middle Class, Not Poor

The terminology here can be confusing. When progressives talk about the working class, typically they mean the poor. But the poor, in the bottom 30% of American families, are very different from Americans who are literally in the middle: the middle 50% of families whose median income was $64,000 in 2008. That is the true “middle class,” and they call themselves either “middle class” or “working class.”

“The thing that really gets me is that Democrats try to offer policies (paid sick leave! minimum wage!) that would *help* the

working class," a friend just wrote me. A few days' paid leave ain't gonna support a family. Neither is minimum wage. WWC men aren't interested in working at McDonald's for $15 per hour instead of $9.50. What they want is what my father-in-law had: steady, stable, full-time jobs that deliver a solid middle-class life to the 75% of Americans who don't have a college degree. Trump promises that. I doubt he'll deliver, but at least he understands what they need.

Understand Working-Class Resentment of the Poor

Remember when President Obama sold Obamacare by pointing out that it delivered health care to 20 million people? Just another program that taxed the middle class to help the poor, said the WWC, and in some cases that's proved true: The poor got health insurance while some Americans just a notch richer saw their premiums rise.

Progressives have lavished attention on the poor for over a century. That (combined with other factors) led to social programs targeting them. Means-tested programs that help the poor but exclude the middle may keep costs and tax rates lower, but they are a recipe for class conflict. Example: 28.3% of poor families receive child-care subsidies, which are largely nonexistent for the middle class. So my sister-in-law worked full-time for Head Start, providing free child care for poor women while earning so little that she almost couldn't pay for

her own. She resented this, especially the fact that some of the kids' moms did not work. One arrived late one day to pick up her child, carrying shopping bags from Macy's. My sister-in-law was livid.

J.D. Vance's much-heralded *Hillbilly Elegy* captures this resentment. Hard-living families like that of Vance's mother live alongside settled families like that of his biological father. While the hard-living succumb to despair, drugs, or alcohol, settled families keep to the straight and narrow, like my parents-in-law, who owned their home and sent both sons to college. To accomplish that, they lived a life of rigorous thrift and self-discipline. Vance's book passes harsh judgment on his hard-living relatives, which is not uncommon among settled families who kept their nose clean through sheer force of will. This is a second source of resentment against the poor.

Other books that get at this are *Hard Living on Clay Street* (1972) and *Working-Class Heroes* (2003).

Understand How Class Divisions Have Translated into Geography

The best advice I've seen so far for Democrats is the recommendation that hipsters move to Iowa. Class conflict now closely tracks the urban-rural divide. In the huge red plains between the thin blue coasts, shockingly high numbers

of working-class men are unemployed or on disability, fueling a wave of despair deaths in the form of the opioid epidemic.

Vast rural areas are withering away, leaving trails of pain. When did you hear any American politician talk about that? Never.

Jennifer Sherman's *Those Who Work, Those Who Don't* (2009) covers this well.

If You Want to Connect with White Working-Class Voters, Place Economics at the Center

"The white working class is just so stupid. Don't they realize Republicans just use them every four years, and then screw them?" I have heard some version of this over and over again, and it's actually a sentiment the WWC agrees with, which is why they rejected the Republican establishment this year. But to them, the Democrats are no better.

Both parties have supported free-trade deals because of the net positive GDP gains, overlooking the blue-collar workers who lost work as jobs left for Mexico or Vietnam. These are precisely the voters in the crucial swing states of Ohio, Michigan, and Pennsylvania that Democrats have so long ignored. Excuse me. Who's stupid?

One key message is that trade deals are far more expensive than we've treated them, because sustained job development and training programs need to be counted as part of their costs.

At a deeper level, both parties need an economic program that can deliver middle-class jobs. Republicans have one: Unleash American business. Democrats? They remain obsessed with cultural issues. I fully understand why transgender bathrooms are important, but I also understand why progressives' obsession with prioritizing cultural issues infuriates many Americans whose chief concerns are economic.

Back when blue-collar voters used to be solidly Democratic (1930–1970), good jobs were at the core of the progressive agenda. A modern industrial policy would follow Germany's path. (Want really good scissors? Buy German.) Massive funding is needed for community college programs linked with local businesses to train workers for well-paying new economy jobs. Clinton mentioned this approach, along with 600,000 other policy suggestions. She did not stress it.

Avoid the Temptation to Write Off Blue-Collar Resentment as Racism

Economic resentment has fueled racial anxiety that, in some Trump supporters (and Trump himself), bleeds into open racism. But to write off WWC anger as nothing more than racism is intellectual comfort food, and it is dangerous.

National debates about policing are fueling class tensions today in precisely the same way they did in the 1970s, when college kids derided policemen as "pigs." This is a recipe for

class conflict. Being in the police is one of the few good jobs open to Americans without a college education. Police get solid wages, great benefits, and a respected place in their communities. For elites to write them off as racists is a telling example of how, although race- and sex-based insults are no longer acceptable in polite society, class-based insults still are.

I do not defend police who kill citizens for selling cigarettes. But the current demonization of the police underestimates the difficulty of ending police violence against communities of color. Police need to make split-second decisions in life-threatening situations. I don't. If I had to, I might make some poor decisions too.

Saying this is so unpopular that I risk making myself a pariah among my friends on the left coast. But the biggest risk today for me and other Americans is continued class cluelessness. If we don't take steps to bridge the class culture gap, when Trump proves unable to bring steel back to Youngstown, Ohio, the consequences could turn dangerous.

In 2010, while on a book tour for *Reshaping the Work-Family Debate*, I gave a talk about all of this at the Harvard Kennedy School. The woman who ran the speaker series, a major Democratic operative, liked my talk. "You are saying exactly what the Democrats need to hear," she mused, "and they'll never listen." I hope now they will.

Originally published in November 2016. Reprint H03913

The Truth About Blockchain

by Marco Iansiti and Karim R. Lakhani

CONTRACTS, TRANSACTIONS, AND THE RECORDS of them are among the defining structures in our economic, legal, and political systems. They protect assets and set organizational boundaries. They establish and verify identities and chronicle events. They govern interactions among nations, organizations, communities, and individuals. They guide managerial and social action. And yet these critical tools and the bureaucracies formed to manage them have not kept up with the economy's digital transformation. They're like a rush-hour gridlock trapping a Formula 1 race car. In a digital world, the way we regulate and maintain administrative control has to change.

Blockchain promises to solve this problem. The technology at the heart of bitcoin and other virtual currencies, blockchain is an open, distributed ledger that can record transactions between two parties efficiently and in a verifiable and permanent way. The ledger itself can also be programmed to

trigger transactions automatically. (See the sidebar “How Blockchain Works.”)

With blockchain, we can imagine a world in which contracts are embedded in digital code and stored in transparent, shared databases, where they are protected from deletion, tampering, and revision. In this world every agreement, every process, every task, and every payment would have a digital record and signature that could be identified, validated, stored, and shared. Intermediaries like lawyers, brokers, and bankers might no longer be necessary. Individuals, organizations, machines, and algorithms would freely transact and interact with one another with little friction. This is the immense potential of blockchain.

How Blockchain Works

Here are five basic principles underlying the technology.

1. **Distributed database.** Each party on a blockchain has access to the entire database and its complete history. No single party controls the data or the information. Every party can verify the records of its transaction partners directly, without an intermediary.
2. **Peer-to-peer transmission.** Communication occurs directly between peers instead of through a central node. Each node stores and forwards information to all other nodes.

3. **Transparency with pseudonymity.** Every transaction and its associated value are visible to anyone with access to the system. Each node, or user, on a blockchain has a unique 30-plus-character alphanumeric address that identifies it. Users can choose to remain anonymous or provide proof of their identity to others. Transactions occur between blockchain addresses.

4. **Irreversibility of records.** Once a transaction is entered in the database and the accounts are updated, the records cannot be altered, because they're linked to every transaction record that came before them (hence the term "chain"). Various computational algorithms and approaches are deployed to ensure that the recording on the database is permanent, chronologically ordered, and available to all others on the network.

5. **Computational logic.** The digital nature of the ledger means that blockchain transactions can be tied to computational logic and in essence programmed. So users can set up algorithms and rules that automatically trigger transactions between nodes.

Indeed, virtually everyone has heard the claim that blockchain will revolutionize business and redefine companies and economies. Although we share the enthusiasm for its potential, we worry about the hype. It's not just security issues (such as the 2014 collapse of one bitcoin exchange and the more recent hacks of others) that concern us. Our experience studying technological innovation tells us that if there's to be a

blockchain revolution, many barriers—technological, governance, organizational, and even societal—will have to fall. It would be a mistake to rush headlong into blockchain innovation without understanding how it is likely to take hold.

Idea in Brief

The Hype

We've all heard that blockchain will revolutionize business, but it's going to take a lot longer than many people claim.

The Reason

Like TCP/IP (on which the internet was built), blockchain is a foundational technology that will require broad coordination. The level of complexity—technological, regulatory, and social—will be unprecedented.

The Truth

The adoption of TCP/IP suggests blockchain will follow a fairly predictable path. While the journey will take years, it's not too early for businesses to start planning.

True blockchain-led transformation of business and government, we believe, is still many years away. That's because blockchain is not a "disruptive" technology, which can attack a traditional business model with a lower-cost solution and overtake incumbent firms quickly. Blockchain is a *foundational* technology: It has the potential to create new

foundations for our economic and social systems. But while the impact will be enormous, it will take decades for blockchain to seep into our economic and social infrastructure. The process of adoption will be gradual and steady, not sudden, as waves of technological and institutional change gain momentum. That insight and its strategic implications are what we'll explore in this article.

Patterns of Technology Adoption

Before jumping into blockchain strategy and investment, let's reflect on what we know about technology adoption and, in particular, the transformation process typical of other foundational technologies. One of the most relevant examples is distributed computer networking technology, seen in the adoption of TCP/IP (transmission control protocol/internet protocol), which laid the groundwork for the development of the internet.

Introduced in 1972, TCP/IP first gained traction in a *single-use* case: as the basis for e-mail among the researchers on ARPAnet, the U.S. Department of Defense precursor to the commercial internet. Before TCP/IP, telecommunications architecture was based on "circuit switching," in which connections between two parties or machines had to be preestablished and sustained throughout an exchange. To ensure that any two nodes could communicate, telecom service providers and equipment

manufacturers had invested billions in building dedicated lines.

TCP/IP turned that model on its head. The new protocol transmitted information by digitizing it and breaking it up into very small packets, each including address information. Once released into the network, the packets could take any route to the recipient. Smart sending and receiving nodes at the network's edges could disassemble and reassemble the packets and interpret the encoded data. There was no need for dedicated private lines or massive infrastructure. TCP/IP created an open, shared public network without any central authority or party responsible for its maintenance and improvement.

Traditional telecommunications and computing sectors looked on TCP/IP with skepticism. Few imagined that robust data, messaging, voice, and video connections could be established on the new architecture or that the associated system could be secure and scale up. But during the late 1980s and 1990s, a growing number of firms, such as Sun, NeXT, Hewlett-Packard, and Silicon Graphics, used TCP/IP, in part to create *localized* private networks within organizations. To do so, they developed building blocks and tools that broadened its use beyond e-mail, gradually replacing more-traditional local network technologies and standards. As organizations adopted these building blocks and tools, they saw dramatic gains in productivity.

TCP/IP burst into broad public use with the advent of the World Wide Web in the mid-1990s. New technology companies quickly emerged to provide the "plumbing"—the hardware, software, and services needed to connect to the now-public network and exchange information. Netscape commercialized browsers, web servers, and other tools and components that aided the development and adoption of internet services and applications. Sun drove the development of Java, the application-programming language. As information on the web grew exponentially, Infoseek, Excite, Alta-Vista, and Yahoo were born to guide users around it.

Once this basic infrastructure gained critical mass, a new generation of companies took advantage of low-cost connectivity by creating internet services that were compelling *substitutes* for existing businesses. CNET moved news online. Amazon offered more books for sale than any bookshop. Priceline and Expedia made it easier to buy airline tickets and brought unprecedented transparency to the process. The ability of these newcomers to get extensive reach at relatively low cost put significant pressure on traditional businesses like newspapers and brick-and-mortar retailers.

Relying on broad internet connectivity, the next wave of companies created novel, *transformative* applications that fundamentally changed the way businesses created and captured value. These companies were built on a new peer-to-peer architecture and generated value by coordinating

distributed networks of users. Think of how eBay changed online retail through auctions, Napster changed the music industry, Skype changed telecommunications, and Google, which exploited user-generated links to provide more relevant results, changed web search.

Ultimately, it took more than 30 years for TCP/IP to move through all the phases—single use, localized use, substitution, and transformation—and reshape the economy. Today more than half the world's most valuable public companies have internet-driven, platform-based business models. The very foundations of our economy have changed. Physical scale and unique intellectual property no longer confer unbeatable advantages; increasingly, the economic leaders are enterprises that act as "keystones," proactively organizing, influencing, and coordinating widespread networks of communities, users, and organizations.

The New Architecture

Blockchain—a peer-to-peer network that sits on top of the internet—was introduced in October 2008 as part of a proposal for bitcoin, a virtual currency system that eschewed a central authority for issuing currency, transferring ownership, and confirming transactions. Bitcoin is the first application of blockchain technology.

The parallels between blockchain and TCP/IP are clear. Just as e-mail enabled bilateral messaging, bitcoin enables bilateral financial transactions. The development and maintenance of blockchain is open, distributed, and shared—just like TCP/IP's. A team of volunteers around the world maintains the core software. And just like e-mail, bitcoin first caught on with an enthusiastic but relatively small community.

TCP/IP unlocked new economic value by dramatically lowering the cost of connections. Similarly, blockchain could dramatically reduce the cost of transactions. It has the potential to become the system of record for all transactions. If that happens, the economy will once again undergo a radical shift, as new, blockchain-based sources of influence and control emerge.

Consider how business works now. Keeping ongoing records of transactions is a core function of any business. Those records track past actions and performance and guide planning for the future. They provide a view not only of how the organization works internally but also of the organization's outside relationships. Every organization keeps its own records, and they're private. Many organizations have no master ledger of all their activities; instead records are distributed across internal units and functions. The problem is, reconciling transactions across individual and private ledgers takes a lot of time and is prone to error.

For example, a typical stock transaction can be executed within microseconds, often without human intervention. However, the settlement—the ownership transfer of the stock—can take as long as a week. That's because the parties have no access to each other's ledgers and can't automatically verify that the assets are in fact owned and can be transferred. Instead a series of intermediaries act as guarantors of assets as the record of the transaction traverses organizations and the ledgers are individually updated.

In a blockchain system, the ledger is replicated in a large number of identical databases, each hosted and maintained by an interested party. When changes are entered in one copy, all the other copies are simultaneously updated. So as transactions occur, records of the value and assets exchanged are permanently entered in all ledgers. There is no need for third-party intermediaries to verify or transfer ownership. If a stock transaction took place on a blockchain-based system, it would be settled within seconds, securely and verifiably. (The infamous hacks that have hit bitcoin exchanges exposed weaknesses not in the blockchain itself but in separate systems linked to parties using the blockchain.)

A Framework for Blockchain Adoption

If bitcoin is like early e-mail, is blockchain decades from reaching its full potential? In our view the answer is a qualified

yes. We can't predict exactly how many years the transformation will take, but we can guess which kinds of applications will gain traction first and how blockchain's broad acceptance will eventually come about.

In our analysis, history suggests that two dimensions affect how a foundational technology and its business use cases evolve. The first is novelty—the degree to which an application is new to the world. The more novel it is, the more effort will be required to ensure that users understand what problems it solves. The second dimension is complexity, represented by the level of ecosystem coordination involved—the number and diversity of parties that need to work together to produce value with the technology. For example, a social network with just one member is of little use; a social network is worthwhile only when many of your own connections have signed on to it. Other users of the application must be brought on board to generate value for all participants. The same will be true for many blockchain applications. And, as the scale and impact of those applications increase, their adoption will require significant institutional change.

How foundational technologies take hold

The adoption of foundational technologies typically happens in four phases. Each phase is defined by the novelty of the applications and the complexity of the coordination efforts needed to make them workable. Applications low in novelty and complexity gain acceptance first.

Applications high in novelty and complexity take decades to evolve but can transform the economy. TCP/IP technology, introduced on ARPAnet in 1972, has already reached the transformation phase, but blockchain applications (in white) are in their early days.

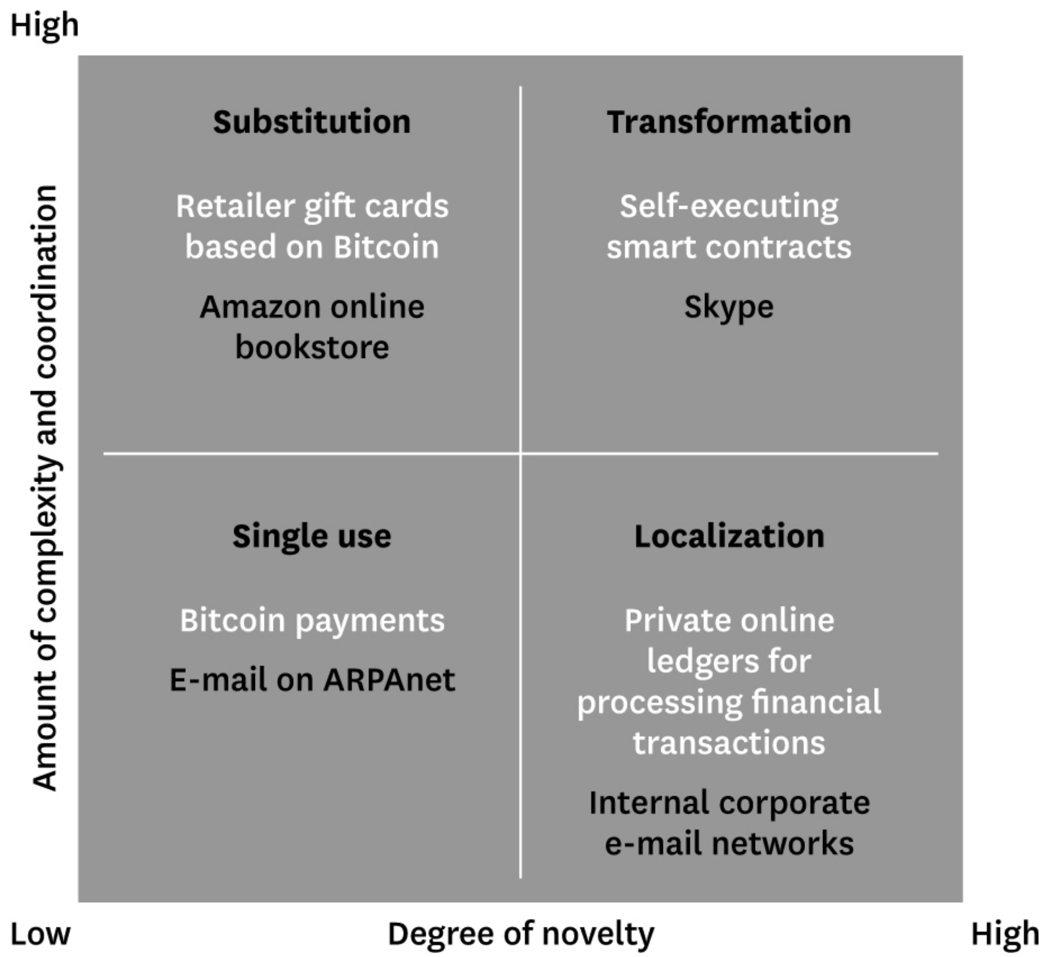

We've developed a framework that maps innovations against these two contextual dimensions, dividing them into quadrants. (See the exhibit "How foundational technologies take hold.") Each quadrant represents a stage of technology development. Identifying which one a blockchain innovation

falls into will help executives understand the types of challenges it presents, the level of collaboration and consensus it needs, and the legislative and regulatory efforts it will require. The map will also suggest what kind of processes and infrastructure must be established to facilitate the innovation's adoption. Managers can use it to assess the state of blockchain development in any industry, as well as to evaluate strategic investments in their own blockchain capabilities.

Single use

In the first quadrant are low-novelty and low-coordination applications that create better, less costly, highly focused solutions. E-mail, a cheap alternative to phone calls, faxes, and snail mail, was a single-use application for TCP/IP (even though its value rose with the number of users). Bitcoin, too, falls into this quadrant. Even in its early days, bitcoin offered immediate value to the few people who used it simply as an alternative payment method. (You can think of it as a complex e-mail that transfers not just information but also actual value.) At the end of 2016 the value of bitcoin transactions was expected to hit $92 billion. That's still a rounding error compared with the $411 trillion in total global payments, but bitcoin is growing fast and increasingly important in contexts such as instant payments and foreign currency and asset trading, where the present financial system has limitations.

Localization

The second quadrant comprises innovations that are relatively high in novelty but need only a limited number of users to create immediate value, so it's still relatively easy to promote their adoption. If blockchain follows the path network technologies took in business, we can expect blockchain innovations to build on single-use applications to create local private networks on which multiple organizations are connected through a distributed ledger.

Much of the initial private blockchain-based development is taking place in the financial services sector, often within small networks of firms, so the coordination requirements are relatively modest. Nasdaq is working with Chain.com, one of many blockchain infrastructure providers, to offer technology for processing and validating financial transactions. Bank of America, JPMorgan, the New York Stock Exchange, Fidelity Investments, and Standard Chartered are testing blockchain technology as a replacement for paper-based and manual transaction processing in such areas as trade finance, foreign exchange, cross-border settlement, and securities settlement. The Bank of Canada is testing a digital currency called CAD-coin for interbank transfers. We anticipate a proliferation of private blockchains that serve specific purposes for various industries.

Substitution

The third quadrant contains applications that are relatively low in novelty because they build on existing single-use and localized applications, but are high in coordination needs because they involve broader and increasingly public uses. These innovations aim to replace entire ways of doing business. They face high barriers to adoption, however; not only do they require more coordination but the processes they hope to replace may be full-blown and deeply embedded within organizations and institutions. Examples of substitutes include cryptocurrencies—new, fully formed currency systems that have grown out of the simple bitcoin payment technology. The critical difference is that a cryptocurrency requires every party that does monetary transactions to adopt it, challenging governments and institutions that have long handled and overseen such transactions. Consumers also have to change their behavior and understand how to implement the new functional capability of the cryptocurrency.

A recent experiment at MIT highlights the challenges ahead for digital currency systems. In 2014 the MIT Bitcoin Club provided each of MIT's 4,494 undergraduates with $100 in bitcoin. Interestingly, 30% of the students did not even sign up for the free money, and 20% of the sign-ups converted the bitcoin to cash within a few weeks. Even the technically savvy had a tough time understanding how or where to use bitcoin.

One of the most ambitious substitute blockchain applications is Stellar, a nonprofit that aims to bring affordable financial

services, including banking, micropayments, and remittances, to people who've never had access to them. Stellar offers its own virtual currency, lumens, and also allows users to retain on its system a range of assets, including other currencies, telephone minutes, and data credits. Stellar initially focused on Africa, particularly Nigeria, the largest economy there. It has seen significant adoption among its target population and proved its cost-effectiveness. But its future is by no means certain, because the ecosystem coordination challenges are high. Although grassroots adoption has demonstrated the viability of Stellar, to become a banking standard, it will need to influence government policy and persuade central banks and large organizations to use it. That could take years of concerted effort.

Transformation

Into the last quadrant fall completely novel applications that, if successful, could change the very nature of economic, social, and political systems. They involve coordinating the activity of many actors and gaining institutional agreement on standards and processes. Their adoption will require major social, legal, and political change.

"Smart contracts" may be the most transformative blockchain application at the moment. These automate payments and the transfer of currency or other assets as negotiated conditions are met. For example, a smart contract

might send a payment to a supplier as soon as a shipment is delivered. A firm could signal via blockchain that a particular good has been received—or the product could have GPS functionality, which would automatically log a location update that, in turn, triggered a payment. We've already seen a few early experiments with such self-executing contracts in the areas of venture funding, banking, and digital rights management.

The implications are fascinating. Firms are built on contracts, from incorporation to buyer-supplier relationships to employee relations. If contracts are automated, then what will happen to traditional firm structures, processes, and intermediaries like lawyers and accountants? And what about managers? Their roles would all radically change. Before we get too excited here, though, let's remember that we are decades away from the widespread adoption of smart contracts. They cannot be effective, for instance, without institutional buy-in. A tremendous degree of coordination and clarity on how smart contracts are designed, verified, implemented, and enforced will be required. We believe the institutions responsible for those daunting tasks will take a long time to evolve. And the technology challenges—especially security—are daunting.

Guiding Your Approach to Blockchain Investment

How should executives think about blockchain for their own organizations? Our framework can help companies identify the right opportunities.

For most, the easiest place to start is single-use applications, which minimize risk because they aren't new and involve little coordination with third parties. One strategy is to add bitcoin as a payment mechanism. The infrastructure and market for bitcoin are already well developed, and adopting the virtual currency will force a variety of functions, including IT, finance, accounting, sales, and marketing, to build blockchain capabilities. Another low-risk approach is to use blockchain internally as a database for applications like managing physical and digital assets, recording internal transactions, and verifying identities. This may be an especially useful solution for companies struggling to reconcile multiple internal databases. Testing out single-use applications will help organizations develop the skills they need for more-advanced applications. And thanks to the emergence of cloud-based blockchain services from both start-ups and large platforms like Amazon and Microsoft, experimentation is getting easier all the time.

Localized applications are a natural next step for companies. We're seeing a lot of investment in private blockchain networks right now, and the projects involved seem poised for real short-term impact. Financial services companies, for example, are finding that the private blockchain networks they've set up

with a limited number of trusted counterparties can significantly reduce transaction costs.

Organizations can also tackle specific problems in transactions across boundaries with localized applications. Companies are already using blockchain to track items through complex supply chains, for instance. This is happening in the diamond industry, where gems are being traced from mines to consumers. The technology for such experiments is now available off-the-shelf.

Developing substitute applications requires careful planning, since existing solutions may be difficult to dislodge. One way to go may be to focus on replacements that won't require end users to change their behavior much but present alternatives to expensive or unattractive solutions. To get traction, substitutes must deliver functionality as good as a traditional solution's and must be easy for the ecosystem to absorb and adopt. First Data's foray into blockchain-based gift cards is a good example of a well-considered substitute. Retailers that offer them to consumers can dramatically lower costs per transaction and enhance security by using blockchain to track the flows of currency within accounts—without relying on external payment processors. These new gift cards even allow transfers of balances and transaction capability between merchants via the common ledger.

Transformative applications are still far away. But it makes sense to evaluate their possibilities now and invest in

developing technology that can enable them. They will be most powerful when tied to a new business model in which the logic of value creation and capture departs from existing approaches. Such business models are hard to adopt but can unlock future growth for companies.

Consider how law firms will have to change to make smart contracts viable. They'll need to develop new expertise in software and blockchain programming. They'll probably also have to rethink their hourly payment model and entertain the idea of charging transaction or hosting fees for contracts, to name just two possible approaches. Whatever tack they take, executives must be sure they understand and have tested the business model implications before making any switch.

Transformative scenarios will take off last, but they will also deliver enormous value. Two areas where they could have a profound impact: large-scale public identity systems for such functions as passport control, and algorithm-driven decision making in the prevention of money laundering and in complex financial transactions that involve many parties. We expect these applications won't reach broad adoption and critical mass for at least another decade and probably more.

Transformative applications will also give rise to new platform-level players that will coordinate and govern the new ecosystems. These will be the Googles and Facebooks of the next generation. It will require patience to realize such opportunities. Though it may be premature to start making

significant investments in them now, developing the required foundations for them—tools and standards—is still worthwhile.

In addition to providing a good template for blockchain's adoption, TCP/IP has most likely smoothed the way for it. TCP/IP has become ubiquitous, and blockchain applications are being built on top of the digital data, communication, and computation infrastructure, which lowers the cost of experimentation and will allow new use cases to emerge rapidly.

With our framework, executives can figure out where to start building their organizational capabilities for blockchain today. They need to ensure that their staffs learn about blockchain, to develop company-specific applications across the quadrants we've identified, and to invest in blockchain infrastructure.

But given the time horizons, barriers to adoption, and sheer complexity involved in getting to TCP/IP levels of acceptance, executives should think carefully about the risks involved in experimenting with blockchain. Clearly, starting small is a good way to develop the know-how to think bigger. But the level of investment should depend on the context of the company and the industry. Financial services companies are already well down the road to blockchain adoption. Manufacturing is not.

No matter what the context, there's a strong possibility that blockchain will affect your business. The very big question is

when.

Further Reading

TO LEARN MORE ABOUT technology adoption, go to these articles on HBR.org:

- **"Digital Ubiquity: How Connections, Sensors, and Data Are Revolutionizing Business,"** Marco Iansiti and Karim R. Lakhani
- **"Strategy as Ecology,"** Marco Iansiti and Roy Levien
- **"Right Tech, Wrong Time,"** Ron Adner and Rahul Kapoor

Originally published in January–February 2017. Reprint R1701J

The Edison of Medicine

by Steven Prokesch

ONE MORNING LAST YEAR, James Dahlman came to Bob Langer's office at MIT's Koch Institute for Integrative Cancer Research to say good-bye. He was meeting with Langer and Dan Anderson —his doctoral advisers. The 29-year-old was about to take up his first faculty position, in the biomedical engineering department at Georgia Tech, and he wanted their advice.

"Do something that's big," Langer told him. "Do something that really can change the world rather than something incremental."

These were not just inspirational words for a former student. They are the watchcry that has guided Langer, a chemical engineer and a pioneer in the fields of controlled-release drug delivery and tissue engineering, throughout his four-decade career at MIT. And they are part of the formula that has made Langer Lab one of the most productive research facilities in the world.

Academic, corporate, and government labs—indeed, anyone leading a group of highly talented people from disparate fields—could learn much from Langer's model. He has a five-pronged approach to accelerating the pace of discoveries and ensuring that they make it out of academia and into the real world as products. It includes a focus on high-impact ideas, a process for crossing the proverbial "valley of death" between research and commercial development, methods for facilitating multidisciplinary collaboration, ways to make the constant turnover of researchers and the limited duration of project funding a plus, and a leadership style that balances freedom and support.

The United States alone spends roughly $500 billion a year on research, but "much of that is mundane," says H. Kent Bowen, an emeritus professor at Harvard Business School who has spent years studying academic and corporate labs. "If there were more highly collaborative, Langer-like labs that focused on high-impact research, the United States would realize its enormous potential for creating wealth."

Langer's achievements are remarkable on several counts. His h-index score, a measure of the number of a scholar's published papers and how often they have been cited, is 230—the highest of any engineer ever. His more than 1,100 current and pending patents have been licensed or sublicensed to some 300 pharmaceutical, chemical, biotechnology, and medical device companies, earning him the nickname "the Edison of

medicine." Alone or in collaboration, his lab has given rise to 40 companies, all but one of which are still in existence, either as independent entities or as part of acquiring companies. Collectively, they have an estimated market value of more than $23 billion—excluding Living Proof, a hair products company that Unilever is acquiring for an undisclosed sum.

A final "product" of the lab is people: Scores of the roughly 900 researchers who have earned graduate degrees or worked as postdocs at the lab have gone on to distinguished careers in academia, business, and venture capital. Fourteen have been inducted into the National Academy of Engineering, 12 into the National Academy of Medicine.

The multidisciplinary approach is still a work in progress in academia, but it has been gathering steam there over the past decade or so, reflecting universities' growing interest in tackling real-world problems and spawning new businesses and a recognition that doing so often takes diverse expertise. Although it has long been common in the business world, companies too could improve their results by applying elements of Langer's research-to-product process, thereby creating brand-new offerings and refreshing or reinventing their businesses again and again.

Idea in Brief

The Problem

Early-stage research is expensive, risky, and unpredictable—so corporations shy away from it, leaving many opportunities unexplored.

The Solution

By pursuing research aimed at solving society's major problems, companies can make the world a better place *and* make lots of money.

The Model

MIT's Bob Langer has a proven formula for accelerating the pace of discoveries and getting them into the world as products—and it's one that any organization can draw on.

Focus on High-Impact Problems

One of Langer's mantras when choosing projects is: Consider the potential impact on society, not the money. The idea is that if you create something that makes a major difference, the customers and the money will come. It's a profound departure from the approach of many big companies: If an idea for a product is so radically new that discounted cash flow can't be calculated, they often won't pursue it, or they give up when the research hits an obstacle—as ambitious research almost always does.

To Langer, "impact" means the number of people an invention could help. The life sciences enterprises that have

emerged from his lab have the potential to touch nearly 4.7 billion lives, according to Polaris Partners, a venture capital firm that has financed many of them. For example, one of the lab's products, on the market since 1996, is a wafer that can be implanted in the brain to deliver chemotherapy directly to the site of a glioblastoma. Another, recently handed over to a new company—Sigilon, based in Cambridge, Massachusetts—is a potential cure for type 1 diabetes, developed in concert with researchers at other universities: Encasing beta cells in a polymer, the researchers have shown, can protect them from the body's immune system yet allow them to detect the level of sugar in the blood and release the appropriate amounts of insulin.

With such concrete, ambitious projects on the lab's docket, the customers have indeed come: foundations, companies, scientists in other labs, and government agencies including the National Institutes of Health. Foundations and companies currently fund 63% of the lab's $17.3 million annual budget; they range from the Bill & Melinda Gates Foundation and the Prostate Cancer Foundation to Novo Nordisk and Hoffmann-La Roche. "A key reason we decided to work with Bob was his lab's track record in controlled delivery," says Dan Hartman, the director of integrated development and malaria at the Gates Foundation and the chief liaison between the foundation and the lab. "Bob and his team's creativity and technical expertise cannot be overemphasized."

A second criterion for project selection is fit with the lab's core areas: drug delivery, drug development, tissue engineering, and biomaterials. "Most of what we do is at the interface of materials, biology, and medicine," Langer says.

Third, he asks whether it's realistic to believe that the medical and scientific challenges can be met by applying or expanding existing science, either at his lab alone or in collaboration with others.

This approach defies a long-prevailing view about the research-to-product process—that it is linear and looks like this: *Basic research* (endeavors aimed at expanding knowledge of nature, without thought of practical use) leads to *applied,* or *translational, research* (efforts to solve practical problems), which in turn leads to *commercial development* (turning discoveries into actual processes and products)—all culminating in a *scale-up* to mass production. The paradigm can be traced to Vannevar Bush, the head of the National Defense Research Committee and the U.S. Office of Scientific Research and Development during World War II and a leading proponent of strong government support for basic scientific research.

Since the war, universities have conducted the lion's share of basic research, but corporations have participated too: Think of AT&T, Corning, DuPont, and IBM, to name just a few. In recent decades, though, big companies have come to see it as too expensive and risky: Results are slow and unpredictable, and capturing their value can be difficult. So they have increasingly

turned to academia, sometimes buying or licensing discoveries or investing in or acquiring start-ups that develop them, other times funding academic research or having their scientists in academic labs.

However, the linear paradigm was never universally true. From the mid 19th century onward, great researchers have pushed the frontiers of basic science precisely to solve pressing societal problems. The Princeton political scientist Donald E. Stokes coined a term for the space in which they work: *Pasteur's quadrant,* reflecting Louis Pasteur's pursuit of a fundamental understanding of microbiology in order to combat disease and food spoilage. Other examples include Bell Labs, whose scientists made basic discoveries while improving and extending communications systems, and the U.S. Defense Advanced Research Projects Agency, or DARPA—one of the most successful innovation organizations ever.

Langer Lab resides in Pasteur's quadrant too. Although its researchers devote the bulk of their efforts to applied science and engineering that could solve critical problems, in the process they often push the boundaries of basic science. For example, one of Langer's most important discoveries was a way to release large-molecule drugs in the body via porous polymers at designated doses and times over several years. This involved expanding an area of physics and math known as percolation theory.

With some notable exceptions—Corning's efforts in quantum communications and materials for capturing carbon dioxide, IBM's in cognitive computing and smart cities, Alphabet's in health care and self-driving vehicles—firms today aren't striving to connect early-stage research with major real-world applications. "It's very rare, but I don't think it needs to be," says Gary P. Pisano, a professor at Harvard Business School. "If you solve some of society's big problems, you'll actually make a lot of money."

Susan Hockfield, a professor of neuroscience at the Koch Institute and a former president of MIT, agrees. "There's a lot of appropriate concern and skepticism about the state of corporate R&D," she says. "For example, pharma corporate R&D invests significantly in very early stage, exploratory research. Couldn't they be doing better if they partnered more effectively with nonindustry biologists and engineers? And I just finished service on a commission to review the national labs. I'm astonished by what a brilliant idea they are and by the high quality of their research, but could they be turning more of their discoveries into products for the marketplace?"

How to Innovate Like Langer

CORPORATIONS TYPICALLY SHY AWAY from early-stage research because it is expensive, risky, and unpredictable, making it difficult for the organization conducting it to capture the

benefits. They could revitalize their research operations by taking an alternative approach and adopting some or all of the following principles from Langer Lab.

Pursue use-inspired research. Companies could direct their research efforts toward concrete problems whose solutions may hold enormous long-term payoffs in terms of the impact on humanity and the ROI. (Bob Langer estimates that venture capitalists have reaped at least a 50% internal rate of return on their investments in the companies he has helped launch.) Those efforts should be a good fit with the company's deep competencies.

Nurture deep scientific and engineering expertise in a handful of areas. This could bring customers flocking for solutions to their most pressing problems.

Manage intellectual property much more aggressively. Companies could benefit from seeking extremely broad, strong patents. And they could license discoveries they don't want to pursue themselves, both to generate income and to ensure that someone pursues them.

Treat the central research organization as a separate entity, liberated from the incremental demands of established business units. In addition, companies could improve their research efforts if they constrained research projects by time, not by creativity.

Staff labs with great—not merely good—scientists and engineers, with an emphasis on making a difference rather than on job stability. Although a number of companies, including Corning, Genentech, Google, IBM, and Novartis, have postdoc

positions and sabbatical programs for professors, the vast majority of researchers even at those firms are long-term employees. Companies could instead give highly talented people two- to five-year contracts, and perhaps a piece of the action if their work succeeds. They should insist on team players with the communication skills, patience, and curiosity to excel in a multidisciplinary context. This approach would give them more flexibility in attracting the range of talent they might need to tackle complex problems.

Establish consistency over time in the funding of, organizational approach to, and independence of advanced research units. This is no easy task; at GE, for example, R&D funding has yo-yoed from one CEO to the next. Success may require a board with a deep understanding of the R&D function and the willingness to push back against an emphasis on quarterly profits.

Ensure robust leadership. This means finding and supporting research directors who are highly respected in their fields and who explicitly see their role as liberating and nurturing the talent around them. Such leaders will have strong networks that can be tapped for recruitment and collaborations; a vision of how the company's expertise can be applied to create major new businesses that are in keeping with corporate strategy; the ability to communicate that vision to secure internal funding and external support; and the goal of making the research organization's value blatantly apparent—ensuring that the unit is seen as the engine of renewal.

Build a Bridge over the Valley of Death

Choosing the right projects to pursue is just the first step, of course; the path to realization can be long and treacherous. Langer has a formula for getting discoveries through the valley of death separating early-stage research and commercial development.

Focus mostly on "platform technologies"—those with multiple applications

Many corporate and academic labs look to solve specific problems without necessarily thinking beyond them. Langer Lab takes a broader view. In addition to creating a wider market, this strategy allows companies to pursue unanticipated applications, says Terry McGuire, a founding partner of Polaris. For example, Momenta, a company launched in 2001 to exploit new methods for understanding and manipulating the structures of sugar molecules, initially set out to sequence heparins in order to treat diseases such as cancer and acute coronary syndrome. However, it realized early on that it could also use the emerging technology to determine the complex structures in Lovenox, an existing multibillion-dollar drug. That work resulted in a biogeneric product for preventing and treating deep vein thrombosis, which generated more than $1 billion in sales during its first year.

Although the lab's researchers often have a use in mind, sometimes they envision a variety of applications. For example, Langer got the idea for an implantable microchip that could

release drugs for years and could be controlled outside the body while watching a television show on semiconductors; he imagined that chips could not only be used to deliver drugs but also put into TVs to release scents that would enhance the viewing experience.

Obtain a broad patent

MIT has been a pioneer in patenting and licensing academic discoveries. But Langer has been exceptional in his pursuit of especially strong patents. His goal is to limit, sometimes even block, others from claiming rights to the territory so that companies will be willing to expend the money needed to commercialize a discovery—an investment that must typically cover expensive clinical trials and that greatly exceeds the cost of the research. (Some of Langer's secrets: Use "great lawyers" and have them challenge one another's recommendations; eliminate unnecessary words that could restrict a claim; and clearly describe all the terms and supporting experimental tests to prevent ambiguity if the patent is litigated.)

Publish a seminal article in a prestigious journal

Appearing in a journal such as *Nature* or *Science* validates—and advertises—the soundness and importance of the discovery not just to other academics but also to potential business investors.

Prove the concept in animal studies, and don't push the discovery out of the lab too quickly

The reason is twofold: to boost the odds that the discovery will work and to minimize the chances that commercialization efforts will flounder—a common occurrence in universities and even the corporate world.

One recent example of a project that benefited from a measured timetable involved the use of ultrasound to rapidly deliver a broad class of therapeutics, including small molecules, macromolecule biologics, and nucleic acids, directly to the gastrointestinal tract (they previously had to be injected). Despite promising initial results and the eagerness of one of the lab's scientists to start a company to commercialize the discovery, Langer resisted taking that step just yet. He wanted to keep the lab team intact and to continue to work on the technology—for instance, demonstrating its safety through "chronic treatment" studies in large animals (giving them the treatment, say, daily for a month) and developing new formulations that could further enhance the delivery of the drugs.

This extra research, unfettered by commercial timetables, paid off. Over the next 18 months or so, the lab demonstrated that the technology could deliver a whole new class of drugs (unencapsulated nucleic acids), broadening its potential applications. The team also published more articles on the research in peer-reviewed journals, providing proof that the original data was reliable and replicable. Only then did Langer

agree to help raise funds for a new company, Suono Bio, to take over development.

Reward the researchers

MIT awards inventors one-third of royalty income after expenses and fees. (The rest goes to the researchers' departments or centers, MIT's technology-licensing office, and the university's general fund.) In recent decades a growing number of universities have instituted similar policies, but the approach is still highly unusual in the corporate world.

Involve the researchers in commercial development

Over the years many members of the lab have left for positions at companies that took on their projects, where their passion for getting the technology to market has proved as important as their expertise. "One of the reasons a lot of the companies have done well is that the champions have been our students who've gone to them," Langer says. "They really believed in what they did in the lab and wanted to make it a reality." Other researchers have advised companies while remaining at the lab or after moving on to other universities. Langer himself serves on the boards of 10 Boston-area start-ups that have emerged from his work. While a growing number of universities have relaxed restrictions on professors' involving themselves in commercial ventures and have even encouraged commercialization by launching incubators and accelerators,

there are still mixed feelings about such activities at many places that lack MIT's established entrepreneurial culture. And in the corporate world, it's highly unusual for scientists to become deeply involved in commercialization.

Make licenses contingent on using the technology

If a firm doesn't make use of technology it has licensed from the lab, it can be made to relinquish the license. And consider how the wafer for treating brain tumors came to market: A company uninterested in the treatment happened to buy the firm that had licensed the technology. MIT got it to agree to launch a start-up to develop the wafer in return for a lower licensing fee. Few universities—or companies—manage their patents as aggressively as MIT does. Consequently, many of their potentially useful discoveries aren't exploited.

Forge a Collaborative Multidisciplinary Team

A team working on an oral drug-delivery device that could sit in the stomach gradually releasing medicine for weeks or months came up with a star-shaped design. Then a mechanical engineer with modeling experience joined the effort and began to ask questions. Why had the team chosen a star? Why not other shapes? The team evaluated several possibilities, including hexagons and a variety of stars, and found that a six-pointed star performed best in terms of its ability to fit inside a capsule and stay in the stomach. The new team member also

raised considerations about the stiffness of the arms and center, the strength of the elastomer at the interface, and the size of the unfolded device. This turned the conversation to materials that might enable the device to last longer.

"That's what happens when you bring together folks with different backgrounds," says Giovanni Traverso, a Harvard gastroenterologist, biomedical engineer, and MIT research affiliate who heads the team. "It leads to new insights and new ways of thinking about the problem." The teams at Langer Lab include chemical, mechanical, and electrical engineers; molecular biologists; medical clinicians; veterinarians; materials scientists; physicists; and pharmaceutical chemists. Members from different disciplines sit side by side in the labs and offices that honeycomb the sixth floor of the Koch Institute.

Multidisciplinary labs are sprouting up as academia recognizes their value in tackling challenges ranging from cancer to global warming. (One of the hallmarks of the Stand Up to Cancer campaign is its funding of such teams.) But the revolution is still in early days. The 2016 MIT report "Convergence: The Future of Health," coauthored by Susan Hockfield, highlights the importance of bringing together engineering, physical, computational, mathematical, and biomedical sciences "to help solve many of the world's grand challenges." It calls for ambitious reforms in education, industry, and government, including the creation of a "culture

of convergence" in academia and industry and changes to government research-funding practices.

Langer's reputation, the challenges his lab takes on, and the career opportunities afforded, including the chance to participate in start-ups, attract lots of applicants. The lab has 119 researchers from all over the world, plus 30 to 40 undergraduates each semester. It receives 4,000 to 5,000 applications for the 10 to 20 postdoc positions that open up each year and conducts global searches when specialized skills are needed for particular projects.

It's a given that applicants must have outstanding academic credentials and be highly motivated. Beyond that, the leadership team of Langer, Traverso, and Ana Jaklenec, a biomedical engineer and MIT staff scientist, looks for people who "are nice, get along well with others, and are good communicators"—vital qualities given that the lab's researchers must constantly explain their fields to coworkers and find ways to conduct experiments that work for everyone. Differences in technical languages, work practices, values, and even ways of defining problems constitute one of the most formidable challenges of a multidisciplinary lab, says Hockfield, a champion of convergence during her eight years at MIT's helm.

Jaklenec showed me a whiteboard filled with equations. It was from a meeting of two postdocs—a biologist and a biomedical engineer who were collaborating on a single-

injection polio vaccine that could stay in the body and be released in pulses over time. The biologist was exploring the mechanism that degrades the strain of virus used in the vaccine, while the biomedical engineer was working on thermostabilization. The two encountered a problem: Their data sets didn't make sense together. It turned out that they had run their experiments with different concentrations of the vaccine: The engineer's were those used clinically, while the biologist's were those called for by the analytical methods of her field. The researchers had to align their experiments so that they could compare results. Such issues are not uncommon. "The challenge is to get people to talk the same language and also recognize that for certain things, there's no single expert," Traverso says.

An Unusual Road to High-Impact Research

IN THE EARLY 1970S, AS BOB LANGER was completing a PhD in chemical engineering at MIT, the United States was rocked by the OPEC embargo and the resulting oil crisis—making him a hot commodity in the eyes of oil and chemical companies (he received 20 job offers in the field). An interview at an Exxon operation in Baton Rouge prompted a seminal insight. "One of the engineers said to me, 'If you could just increase the yield of this one chemical by point-one percent, that would be wonderful—that's worth billions of dollars,'" Langer recalls. "I

remember flying back to Boston that night thinking, 'Do I really want to spend my life doing this?'"

He applied to colleges for jobs developing chemistry curricula. When none replied—"probably because as a chemical engineer, I wasn't in the right box"—he wrote to hospitals, "because I wanted to help people." Again he received no offers.

Then a colleague suggested that he contact Judah Folkman, a surgeon at Boston Children's Hospital who had a reputation for hiring unusual people. Folkman had a controversial idea: that cancerous tumors emit chemical signals that stimulate angiogenesis, or the formation of new blood vessels. If the signals could be blocked, Folkman theorized, tumors' growth could be halted. He hired Langer to isolate the first angiogenesis inhibitors. This involved identifying candidates from cartilage, which has no blood supply (Langer got cow bones from a slaughterhouse) and inventing polymer systems that could deliver large molecules over time. Angiogenesis inhibitors ultimately became instrumental in treating a number of cancers, and polymers have become an important way to deliver drugs and vaccines and to help grow new body tissue, including skin, cartilage, and spinal cord.

Langer returned to MIT in 1977 as an assistant professor, initially in the Department of Nutrition and Food Science (because no chemical engineering department at a university would hire him). It gave him tremendous freedom, and he continued working on drug delivery, angiogenesis inhibitors, and tissue engineering, obtaining funding from companies when his ideas proved too radical for government grants. Many senior faculty members of the department didn't

believe in his ideas and suggested that he look for a new job. However, by the mid-1980s his discoveries, publications, and start-ups began winning recognition. One of MIT's 13 Institute Professors, Langer is a member of the National Academies of Sciences, Engineering, and Medicine, and a recipient of the National Medal of Technology and Innovation, the National Medal of Science, the Charles Stark Draper Prize, and the Queen Elizabeth Prize for Engineering.

Even if there is no obvious need or fit for them, Langer often brings in "superstars" who have unusual credentials. "You take a chance on people," he says. "Gio is a good example." Traverso had earned a PhD in molecular biology under Bert Vogelstein, a renowned cancer biologist at Johns Hopkins; his doctoral research involved novel molecular tests for the early detection of colon cancer. When he contacted Langer, he was finishing an internal medicine residency at Boston's Brigham and Women's Hospital and trying to figure out what to do with a gastroenterology fellowship he had landed at Massachusetts General Hospital. He told Langer that although he was interested in developing systems for delivering drugs in the GI tract, he was not an engineer. Langer hired him anyway.

The bet paid off. Traverso demonstrated the concept of several different approaches to delivering drugs through devices in the GI tract. The Gates Foundation saw that the work might solve problems it wanted to address in poor countries and provided significant funding. Grants also came in

from Novo Nordisk (to develop microneedles for internal injections), the Charles Stark Draper Lab (for new ingestible systems), and Hoffmann-La Roche (for the delivery of a new class of drugs).

Embrace Turnover

Like all academic labs, Langer's sees a constant flow of people joining or leaving. Doctoral students typically stay four or five years, postdocs two or three, and undergraduates participate for as little as a semester and as much as four years. Newcomers are perpetually being trained, and people may leave at the peak of their productivity. But Langer and many colleagues think the turnover has positives that vastly outweigh these downsides. Problems are viewed with fresh eyes—he calls it "constant stimulation." The turnover is fairly predictable and tied to the length of projects; even huge grants are structured so that the lab can gradually scale up. The finite tenure of most of the researchers, combined with the limited duration of grants (typically three to five years, with renewals dependent on meeting goals), imposes pressure to get results.

"A lot of cynicism has been thrown on the academic research lab model. We are told it is inefficient," Hockfield says. "But it's brilliant. To bring together people from different generations and levels of experience—it's fantastic. The faculty member has a wealth of experience and understanding and knows the

literature and the history of the field. Students and postdocs have a lot of energy and ambition and crazy ideas. The faculty member helps get those crazy ideas channeled. Undergraduates, wonderfully, often don't know that something's impossible. They don't know enough not to ask unsophisticated questions. There are very few things that make you step back and wonder about your foundational assumptions more than a really smart undergraduate asking, 'Whoa, how does that work?'"

A highly motivated superstar team with limited tenure; an accomplished scientist leader; time-limited projects; intense pressure to get results—it all sounds like the DARPA formula, proof that the model has application far beyond academic settings.

Lead Without Micromanaging

One rainy day at their home on Cape Cod, Langer and his wife, Laura, talked about how his management of the lab differs from the norm. "In my discussions with a range of graduate students at other places, they often describe their research advisers as control freaks—which is understandable, because their lab is their baby," said Laura, who has a PhD in neuroscience from MIT. "They may want to manage every part of the research. It's very hard for them to let their students explore and make mistakes. But not giving people the room to

figure things out themselves can stifle them or train them to not take potentially innovative risks."

Langer nodded in agreement. Under his leadership, everyone is involved in offering ideas for projects and choosing which ones to pursue. "It's a team effort," he said. "It's empowering people; it's letting everybody feel they are valued and that it's OK to suggest things." This stands in contrast to most academic and corporate labs, where the director selects the projects.

Current and former lab members told me that Langer exposes people to possibilities and lets them decide what to work on. Gordana Vunjak-Novakovic, a professor of biomedical engineering and medical sciences at Columbia who worked at the lab in the 1980s and 1990s, says she took that lesson to heart and runs her 40-person lab the same way: "I never tell people what to do but, rather, help them see the possibilities, let them really get excited about one of them, and let them work on their own ideas." Many if not most of Langer's postdocs and research scientists and at least some of the doctoral students are working on several projects. (For a fuller picture of life in Langer Lab, see the profile of two postdocs in the online version of this article, at HBR.org.)

Langer treats Jaklenec and Traverso as coprincipal investigators—another departure from the norm. Power is distributed throughout the lab, accumulated on the basis of people's ideas and initiative and the funding that their research attracts. Langer gives researchers—especially graduate students

—lots of guidance in the beginning, to make sure that they get off to a good start and that projects are optimally structured. He also helps decide which options are considered. For example, at the outset of the project to develop the drug-delivery device that would stay in the stomach for a long period, he and Traverso decided to explore two possibilities: one that would float in the stomach and one that would adhere to the stomach wall. After conducting a feasibility study, they chose to pursue the floating option and figured out what major issues would need to be solved—and then Langer largely bowed out. “After that, I don’t tell people what to do,” he says. “From grade school to high school and college and even to a certain extent graduate school, you’re judged by how well you answer somebody else’s questions. That gives you a grade on a test. But if you think about the way you’re judged in life, I don’t think it is by how good your answers are; it’s by how good your questions are. I want to help people make that transition from giving good answers to asking good questions.”

Real-World Results

SINCE 1987 BOB LANGER and his researchers have helped found 40 companies, often in collaboration with scientists in other labs at MIT and at other institutions. To date all but one have made it. A sampling is below.

Company: Enzytech (acquired by Alkermes)

Year launched: 1987

Products/technology: Microspheres for delivering drugs

Existing or potential applications: Schizophrenia, narcotic addiction, type 2 diabetes

Market capitalization: $7.2 billion (Alkermes)

Company: Moderna

Year launched: 2011

Products/technology: Messenger-RNA-based drugs

Existing or potential applications: Cancer, heart disease, vaccines, infectious diseases, pulmonary disease

Market capitalization: $5 billion

Company: Momenta

Year launched: 2001

Products/technology: Sequencing complex sugar-based therapeutics

Existing or potential applications: Multiple sclerosis and other autoimmune diseases, cardiovascular diseases, cancer

Market capitalization: $840 million

Company: Advanced Inhalation Research (acquired by Acorda)

Year launched: 1997

Products/technology: Drug-delivering aerosols that rely on large particles, which resist clumping

Existing or potential applications: Diabetes, asthma, Parkinson's disease

Market capitalization: $525 million

Company: Selecta

Year launched: 2007

Products/technology: Targeted nanoparticle-based immunotherapies and vaccines

Existing or potential applications: Gout, genetic disorders, allergies, autoimmune diseases, HPV-associated cancers, nicotine addiction, malaria

Market capitalization: $228 million

Sources: Robert Langer, Polaris Partners, public information.
Note: Market capitalizations are as of mid-September 2016 or acquisition date. The value of private companies is based on VC financing.

Gary Pisano sees this philosophy as key to the lab's success. "The tendency would be to say, 'I'm going to tell you what to do so that you can do better and the lab will do better,'" he explains. "But if you do that, you create a different place—people are going to say, 'OK, Bob, you tell me what to do.' He doesn't want that kind of lab. His lab is one where people solve their own problems, and that's why they wind up being great professors and scientists in the business world."

At the same time, Langer makes sure that researchers know they can count on him and on the people in his network if they run into trouble—an approach that Aimee L. Hamilton, an assistant professor of management at the University of Denver who has studied Langer Lab, calls "guided autonomy." His responsiveness is legendary. His iPad seems glued to him, and he uses it to answer e-mails within minutes. Cato T. Laurencin,

a University Professor at the University of Connecticut who earned his PhD under Langer in the 1980s, recalls that a student of his once dug up Langer's cell phone number and called him with a question about a paper Langer had written. "He called her back from Finland 10 minutes later."

Langer also goes out of his way to help people leaving his lab get good jobs, and he stays in touch with hundreds of alumni, providing assistance if needed. (In his farewell meeting with James Dahlman, he offered to go over Dahlman's grant applications.) He is deeply connected to those in his network. For instance, he refers to many of the venture capitalists who have financed his start-ups—a group including Terry McGuire, of Polaris; Noubar Afeyan, of Flagship; and Mark Levin, of Third Rock—as friends, and means it. (Langer, McGuire, and their two daughters vacationed together last year in Bordeaux, and Langer's daughter was in the wedding of McGuire's.)

The investment in his network pays valuable dividends in the form of productive research collaborations, referrals of extraordinary students to his lab, and manpower for the start-ups. Langer not only paves the way for lab members to launch start-ups but also taps his network if a need at one emerges down the road. "Bob often has a great idea of somebody who would be a great fit," says Amy Schulman, the CEO or executive chair of three companies that grew out of Langer Lab. "And people often reach out to Bob when they're thinking of

changing jobs, because he is incredibly discreet and knows a lot of opportunities. So it goes both ways."

When people who have worked with Bob Langer talk about him, one hears a common refrain: He is an integral part of his research-to-product model and a brilliant individual who can't be replicated. But this doesn't mean that his model, including his "Mr. Nice Guy" leadership style, can't be replicated. What if corporations structured their labs like his? What if they nurtured deep expertise in a handful of areas so that customers would come to them with their most pressing problems? What if they enticed superstar researchers by offering opportunities to work on issues that could change the world?

"Maybe companies could set up a research operation where the best of the best are flowing through, trying to do something audacious in a few years rather than spending 30 years there worrying about their next promotion," Gary Pisano says. His Harvard colleague Willy Shih adds that such an approach would not only allow companies to tackle more-ambitious projects but also help them kill mediocre or poor projects faster. "The flow of people through the lab would have the natural consequence of sunsetting ideas that don't stand the test of a fresh look," he points out.

Bob Langer says, "I want to address problems that can change the world and make it a better place. That's the thread

throughout the science I've done my whole life. The companies I've helped found seem like a natural extension. I wanted to see what I did get out to the world; that made a difference to me." By drawing on the Langer Lab values and model, companies could make the world a better place and make lots of money in the process.

Further Reading

FOR MORE ON REVITALIZING your research operations, see these articles on HBR.org.

- **"Getting Your Stars to Collaborate,"** Heidi K. Gardner (January–February 2017)
- **"You Need an Innovation Strategy,"** Gary P. Pisano (June 2015)
- **"'Special Forces' Innovation: How DARPA Attacks Problems,"** Regina E. Dugan and Kaigham J. Gabriel (October 2013)
- **"Rebuilding the R&D Engine in Big Pharma,"** Jean-Pierre Garnier (May 2008)

Originally published in March–April 2017. Reprint R1702L

About the Contributors

RON ADNER is a professor of strategy and entrepreneurship at Dartmouth College's Tuck School of Business, where he holds the David T. McLaughlin D'54, T'55 endowed chair.

SCOTT BERINATO is a senior editor at HBR and the author of *Good Charts: The HBR Guide to Making Smarter, More Persuasive Data Visualizations* (Harvard Business Review Press, 2016), from which this article is adapted.

TOM BLASER is a managing director at The Greatest Good Group (TGG Group).

PETER CAPPELLI is a professor of management at the Wharton School and the author of several books, including *Will College Pay Off? A Guide to the Most Important Financial Decision You'll Ever Make* (PublicAffairs, 2015).

FRANK DOBBIN is a professor of sociology at Harvard University.

LINNEA GANDHI is a managing director at TGG Group.

FRANCESCA GINO is a professor at Harvard Business School, a faculty affiliate of the Behavioral Insights Group at Harvard Kennedy School, and the author of *Sidetracked: Why Our Decisions Get Derailed, and How We Can Stick to the Plan* (Harvard

Business Review Press, 2013). She cochairs an HBS executive education program on applying behavioral economics to organizational problems.

MARCO IANSITI is the David Sarnoff Professor of Business Administration at Harvard Business School.

DANIEL KAHNEMAN is the Eugene Higgins Professor of Psychology Emeritus at Princeton University. He was awarded the Nobel Prize in Economic Sciences in 2002 for his work (with Amos Tversky) on cognitive biases.

ALEXANDRA KALEV is an associate professor of sociology at Tel Aviv University.

ROBERT S. KAPLAN is a senior fellow and the Marvin Bower Professor of Leadership Development, Emeritus, at Harvard Business School. He is the coauthor (with Michael E. Porter) of "How to Solve the Cost Crisis in Health Care" (HBR, September 2011).

RAHUL KAPOOR is an associate professor of management at the Wharton School.

A.G. LAFLEY, the recently retired CEO of Procter & Gamble, serves on the board of Snap Inc. He is the coauthor (with Roger L. Martin) of *Playing to Win: How Strategy Really Works* (Harvard Business Review Press, 2013).

KARIM R. LAKHANI is a professor of business administration at Harvard Business School and the principal investigator of the Crowd Innovation Laboratory at the Harvard Institute for Quantitative Social Science.

ROGER L. MARTIN is a former dean of the Rotman School of Management at the University of Toronto. He is the coauthor (with A.G. Lafley) of *Playing to Win: How Strategy Really Works* (Harvard Business Review Press, 2013).

RITA GUNTHER MCGRATH, a professor of management at Columbia Business School, is a globally recognized expert on strategy, innovation, and growth with an emphasis on corporate entrepreneurship.

MICHAEL E. PORTER is a University Professor at Harvard, based at Harvard Business School in Boston. He is the coauthor (with Robert S. Kaplan) of "How to Solve the Cost Crisis in Health Care" (HBR, September 2011).

STEVEN PROKESCH is a senior editor at *Harvard Business Review*.

ANDREW M. ROSENFIELD is the CEO and managing partner of TGG Group.

ANNA TAVIS is the academic director of Columbia's program in human capital management and the Perspectives editor at *People + Strategy*, a journal for HR executives.

JOAN C. WILLIAMS is Distinguished Professor of Law and Founding Director of the Center of WorkLife Law at the University of California, Hastings College of the Law. Her newest book is *White Working Class: Overcoming Class Cluelessness in America* (Harvard Business Review Press, 2017).

Index

HBR'S 10 MUST READS

BONUS ARTICLE

"Now What?"

Managing #MeToo

By Joan C. Williams and Suzanne Lebsock

The definitive management ideas of the year from **Harvard Business Review.**

2019

HBR's 10 Must Reads series is the definitive collection of ideas and best practices for aspiring and experienced leaders alike. These books offer essential reading selected from the pages of *Harvard Business Review* on topics critical to the success of every manager.

Titles include:

HBR's 10 Must Reads 2015
HBR's 10 Must Reads 2016
HBR's 10 Must Reads 2017
HBR's 10 Must Reads 2018
HBR's 10 Must Reads 2019
HBR's 10 Must Reads for New Managers
HBR's 10 Must Reads on Change Management
HBR's 10 Must Reads on Collaboration
HBR's 10 Must Reads on Communication
HBR's 10 Must Reads on Emotional Intelligence
HBR's 10 Must Reads on Entrepreneurship and Startups
HBR's 10 Must Reads on Innovation
HBR's 10 Must Reads on Leadership
HBR's 10 Must Reads on Leadership for Healthcare
HBR's 10 Must Reads on Leadership Lessons from Sports
HBR's 10 Must Reads on Making Smart Decisions
HBR's 10 Must Reads on Managing Across Cultures
HBR's 10 Must Reads on Managing People

HBR's 10 Must Reads on Managing Yourself
HBR's 10 Must Reads on Mental Toughness
HBR's 10 Must Reads on Sales
HBR's 10 Must Reads on Strategic Marketing
HBR's 10 Must Reads on Strategy
HBR's 10 Must Reads on Strategy for Healthcare
HBR's 10 Must Reads on Teams
HBR's 10 Must Reads: The Essentials

HBR'S 10 MUST READS

The definitive
management ideas
of the year from
Harvard Business Review.
2019

HARVARD BUSINESS REVIEW PRESS
Boston, Massachusetts

Copyright

HBR Press Quantity Sales Discounts

Harvard Business Review Press titles are available at significant quantity discounts when purchased in bulk for client gifts, sales promotions, and premiums. Special editions, including books with corporate logos, customized covers, and letters from the company or CEO printed in the front matter, as well as excerpts of existing books, can also be created in large quantities for special needs.

For details and discount information for both print and ebook formats, contact booksales@harvardbusiness.org, tel. 800-988-0886, or www.hbr.org/bulksales.

Copyright 2019 Harvard Business School Publishing Corporation

All rights reserved

No part of this publication may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form, or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the prior permission of the publisher. Requests for permission should be directed to permissions@hbsp.harvard.edu, or mailed to Permissions, Harvard Business School Publishing, 60 Harvard Way, Boston, Massachusetts 02163.

First eBook Edition: Nov 2018

ISBN: 9781633696426
eISBN: 9781633696433

Contents

Editors' Note

It's never easy to whittle down a year's worth of *Harvard Business Review*'s research, ideas, and advice to the few articles gathered in this volume, but this past year was particularly tough. In addition to staple HBR topics such as leadership and strategy, the complex and difficult issues we were turning over in our minds and discussing in boardrooms and on social media also filled the pages of HBR. Recurring themes included machine learning, the place of business in society, and the implications of intersectionality—where harassment and discrimination can affect any one of the multiple layers of our identity. The standout articles of the year covered an array of topics, from integrating cognitive technology with human work to speaking up—whether as a CEO activist or as a manager amid the #MeToo movement. Our authors gave you new lenses through which to view the evolving context in which we work. This collection of articles showcases these and other critical themes from the past year of *Harvard Business Review*.

We've all been working in teams for years. The challenge today is how to manage work and communication when you and everyone you work with are all on a half-dozen other teams too. "**The Overcommitted Organization**" affirms that some standard advice for working on teams still applies while also providing new strategies for managing this growing modern-day dilemma, from mapping overlap to sharing insights across

projects to helping teams maintain progress when key members are yanked for “all hands on deck” emergencies. Authors Mark Mortensen and Heidi K. Gardner conduct research, teach, and consult on collaboration and leadership issues. They have identified several ways in which both team and organizational leaders can reduce the negative aspects of overlap and take advantage of the benefits, including skill sharing across teams, better time management, and opportunities to learn.

MBA students are taught that companies can’t expect to compete on the basis of management competencies—they’re too easy for rivals to copy, so they won’t sustain competitive advantage over time. However, a decade-long research project undertaken by authors Raffaella Sadun, Nicholas Bloom, and John Van Reenen reveals that the conventional wisdom is flawed, raising the question “**Why Do We Undervalue Competent Management?**” In their study of 12,000 organizations the authors found vast differences in how companies execute 18 core management practices, including such basic ones as setting targets, running operations, and grooming talent. Those differences matter: Companies with strong managerial processes do significantly better on high-level metrics such as profitability, growth, and productivity. The authors identify the main challenges hindering the adoption of essential management practices, suggest solutions, and make the case

that senior leaders should focus on operational excellence as a crucial complement to strategy.

To overcome organizational bias, leaders are relying on people analytics to make data-driven decisions and to hire and promote fairly. But some leaders who take this approach say they can't counteract or reverse bias with data: They can't "apply analytics to the challenges of underrepresented groups at work" because "the relevant data sets don't include enough people to produce reliable insights—the sample size, the *n*, is too small." In "'**Numbers Take Us Only So Far**,'" Facebook's global director of diversity, Maxine Williams, explains why data must be paired with qualitative research to give leaders the insights they need to increase diversity at all levels of their organizations. By drawing on industry or sector data, learning what other companies are doing, and deeply examining the experiences of their own employees, companies can advance their goals of improving diversity and inclusion.

CEOs have always lobbied publicly for political or social issues that are good for their business. But this year we saw a significant phenomenon emerge: "**The New CEO Activists**." Taking stands on issues that are not directly related to their business model and their success can hurt sales (or help them) when consumers respond with their wallets. So why take the risk? Duke's Aaron K. Chatterji and Harvard Business School's Michael W. Toffel offer a guide leaders can use in assessing whether to speak out and how, choosing which issues to weigh

in on, and balancing the likelihood of having a positive effect with the possibility of a backlash.

Artificial intelligence and machine learning have generated lots of hype, but what do they mean for you and your business? In "**Artificial Intelligence for the Real World**," Thomas H. Davenport and Rajeev Ronanki encourage readers to look at AI "through the lens of business capabilities rather than technologies." Instead of a transformative approach, the authors advise, companies should take an incremental approach to developing and implementing AI and focus on augmenting rather than replacing human capabilities. They assert that AI can support three important business needs: automating business processes, gaining insight through data analysis, and engaging with customers and employees. Their four-step framework for integrating AI technologies, along with the real-case examples they provide, will allow companies to explore how they might best use cognitive technologies.

For those who work outside the technology realm, the acronyms AI and AR can sound a bit like alphabet soup. We found value in reading the previous piece and "**Why Every Organization Needs an Augmented Reality Strategy**" together, because that can help define what those acronyms are and how they're used. AR—technologies that superimpose digital data and images on physical objects—has familiar entertainment applications, such as Snapchat and Pokémon Go. But AR is now being used in business in far more consequential ways; Michael

E. Porter and James E. Heppelmann assert that it will become the new interface between humans and machines. They define AR, describe its evolving technology and applications, and discuss its importance. The authors provide both a primer for Luddites and an expansive review of the opportunities AR presents, from expected applications such as logistics and design to surprising ones such as allowing HR to tailor training according to an employee's experience or repeated errors.

Whether we're freelancers who have lost access to the security and support of traditional employers or corporate employees logging in from home offices, the way we work has changed. In "**Thriving in the Gig Economy**," the organizational behavior professors Gianpiero Petriglieri, Susan Ashford, and Amy Wrzesniewski report on their study of freelance workers to understand what it takes to be successful in independent work. They found that the most effective independent workers "cultivate four types of connections—to *place, routines, purpose,* and *people*—that help them endure the emotional ups and downs of their work and gain energy and inspiration from their freedom." Addressing these core areas can help you stay motivated, boost your productivity and focus, and ward off feelings of rootlessness and isolation.

As individuals, we're working in new ways, but the context in which we work and our organizations grow—or fail—has changed too. "**Managing Our Hub Economy**" offers a fascinating, forward-looking, and sometimes chilling examination of the

place of business in society. Hub firms such as Alibaba, Apple, and Amazon create real value for users but also concentrate data and power in the hands of a few companies that employ a tiny fraction of the workforce. Harvard Business School professors Marco Iansiti and Karim R. Lakhani argue that the hub economy will continue to spread across additional industries, concentrating power even more. “To remain competitive, companies will need to use their assets and capabilities differently, transform their core businesses, develop new revenue opportunities, and identify areas that can be defended from encroaching hub firms and others rushing in from previously disconnected economic sectors.”

Another new perspective on an old issue is found in “**The Leader’s Guide to Corporate Culture**.” The conventional wisdom has it that leaders are expected to create and change strategy, but culture is ingrained, unchangeable, and “anchored in unspoken behaviors, mindsets, and social patterns.” Not so, say Harvard Business School professor Boris Groysberg and his coauthors. They argue that it is possible to change your company’s culture, but first you must understand how it works. By integrating findings from more than 100 of the most commonly used social and behavioral models, the authors created a framework that will allow you to model the impact of culture on your business and assess its alignment with your strategy. When properly managed, culture can help leaders

achieve change and build organizations that will thrive in even the most trying times.

Most CEOs and boards are hyperfocused on creating wealth for their shareholders. But managing for the good of the stock is not always the same as managing for the good of the company—especially when it leads to a focus on the short term. In "**The Error at the Heart of Corporate Leadership**," Joseph L. Bower and Lynn S. Paine examine the foundations and flaws of agency theory, which views shareholders as the "owners" of a company and is behind the current widespread idea that corporate managers should make shareholder value their primary concern. The authors offer eight propositions to provide a company-centered model that would have at its core the health of the enterprise instead. Their model would return companies' attention to innovation, strategic renewal, and investment in the future.

Where do we go from here? The #MeToo movement and countless reports of sexual harassment in the workplace are transforming how we manage relationships at work. In "**Now What?**" the legal scholar Joan C. Williams and the feminist historian Suzanne Lebsock explore whether this is really the end of a harassment culture. Companies are moving away from quiet settlements with victims and toward firing abusers. But employers must still follow due process and evaluate the credibility of reports. They need clear policies and fair procedures for handling harassment. The authors surprised

themselves with their closing advice: “If you are being sexually harassed, report it. We’re not sure if we would have advised that, in such a blanket and unnuanced way, even a year ago.”

The most important ideas of the year are at your fingertips in this volume. From ideas on managing your team, to issues for your board and senior executives, to harnessing artificial intelligence and augmented reality, to addressing meaty personnel issues such as diversity and harassment, the articles here will help you address the situations you’re facing today and prepare for what lies on the horizon.

—The Editors

The Overcommitted Organization

by Mark Mortensen and Heidi K. Gardner

A SENIOR EXECUTIVE WE'LL CALL Christine is overseeing the launch of Analytix, her company's new cloud-based big-data platform, and she's expected to meet a tight go-live deadline. Until two weeks ago, her team was on track to do that, but it has since fallen seriously behind schedule. Her biggest frustration: Even though nothing has gone wrong with Analytix, her people keep getting pulled into other projects. She hasn't seen her three key engineers for days, because they've been busy fighting fires around a security breach on another team's product. Now she has to explain to the CEO that she can't deliver as promised—at a time when the company badly needs a successful launch.

Christine's story is hardly unique. Across the world, senior managers and team leaders are increasingly frustrated by conflicts arising from what we refer to as multiteaming—having their people assigned to multiple projects simultaneously. But given the significant benefits of multiteaming, it has become a way of organizational life, particularly in knowledge work. It allows groups to share individuals' time and brainpower across

functional and departmental lines. It increases efficiency, too. Few organizations can afford to have their employees focus on just one project at a time and sit idle between tasks. So companies have optimized human capital somewhat as they would machines in factories, spreading expensive resources across teams that don't need 100% of those resources 100% of the time. As a result, they avoid costly downtime during projects' slow periods, and they can bring highly specialized experts in-house to dip in and out of critical projects as needed. Multiteaming also provides important pathways for knowledge transfer and the dissemination of best practices throughout organizations.

As clear and quantifiable as these advantages are, the costs are substantial and need to be managed, as Christine would attest. Organizations open themselves up to the risk of transmitting shocks across teams when shared members link the fates of otherwise independent projects. And teams discover that the constant entrance and exit of members weakens group cohesion and identity, making it harder to build trust and resolve issues. Individual employees pay a big price as well. They often experience stress, fatigue, and burnout as they struggle to manage their time and engagement across projects.

Over the past 15 years, we have studied collaboration in hundreds of teams, in settings as varied as professional services, oil and gas, high tech, and consumer goods. (See the sidebar "About the Research.") By carefully observing people during

various stages of project-driven work, we have learned a tremendous amount about multiteaming. In this article we discuss why it is so prevalent in today's economy, examine the key problems that crop up for organizational and team leaders, and provide recommendations for how to solve them.

About the Research

Over the past 15 years, we've been measuring both the benefits and the trade-offs of multiteaming in areas such as human capital, resource utilization, quality management, and customer satisfaction. We have conducted:

- **In-depth studies** of eight global professional services firms where multiteaming is the norm, including statistical analyses of their staffing databases and personnel records.
- **A survey** of more than 500 midlevel managers in global companies, representing a wide range of industries and professions, to examine trends across organizations and geographies.
- **Ongoing research** at a 5,000-person technology and services company that is trying to optimize multiteaming. So far, this includes more than 50 interviews with team leaders and executives. We're also designing organizational experiments to test best practices and collect data on outcomes such as efficiency, staff burnout, and customer satisfaction.
- **Ongoing research** on agent-based modeling to understand the behavior of large systems of interconnected teams. We are also using

simulations to model multiteaming, with a focus on understanding the relationship between team size, percentage of overlap among teams, and the number of teams each team member is on.

Why This Matters Now

Even though assigning employees to multiple projects at once is not new, the practice is especially widespread today. In a survey of more than 500 managers in global companies, we found that 81% of those working on teams worked on more than one concurrently. Other research places the number even higher—for example, 95% in knowledge-intensive industries.

Why is multiteaming practically ubiquitous? For several reasons.

First, organizations must draw on expertise in multiple disciplines to solve many large, complex problems. Businesses are tackling cybersecurity risks that span departments as diverse as finance, supply chain, and travel. Energy companies are coordinating global megaprojects, including the opening of new deep-sea resource fields. Transportation and logistics firms are tasked with getting resources from point A to point B on time, irrespective of how remote those points are or what is being delivered. Large-scale manufacturing and construction endeavors, such as aircraft and city infrastructure projects, require tight collaboration between those producing the work and the agencies regulating it. In such contexts, organizations

can't rely on generalists to come up with comprehensive, end-to-end solutions. They must combine the contributions of experts with deep knowledge in various domains. (For more on this, see "Getting Your Stars to Collaborate," HBR, January–February 2017.)

Idea in Brief

The Pros

By assigning people to multiple teams at once, organizations make efficient use of time and brainpower. They also do a better job of solving complex problems and sharing knowledge across groups.

The Cons

Competing priorities and other conflicts can make it hard for teams with overlapping membership to stay on track. Group cohesion often suffers. And people who belong to many teams at once may experience burnout, which hurts engagement and performance.

The Fixes

Leaders can mitigate these risks by building trust and familiarity through launches and skills mapping, identifying which groups are most vulnerable to shocks, improving coordination across teams, and carving out more opportunities for learning.

Second, with crowded markets and reduced geographic and industry barriers, organizations now face greater pressure to keep costs down and stretch resources. One client manager in a professional services firm noted, "To be really good stewards of client dollars, we don't want to pay for five weeks of a specialist's time when what we really need is an intense effort from that

person in week five.” That’s why “bench time” between projects and even slow periods during projects have become increasingly rare. The instant people are underutilized, their organizations put them to work on other things. In our research we found that even senior-level managers were flipping among seven or more projects in a single day—and as many as 25 in a given week. Compounding this, technology makes it easier to track downtime—even if it’s just minutes—and assign employees work or loop them into projects during any lulls.

Third, organizational models are moving away from hierarchical, centralized staffing to give employees more choice in their projects and improve talent development, engagement, and retention. Indeed, in the gig economy, individuals have greater control than ever over the work they do (think open-source software programmers). This has made leading teams an even more critical skill. (For more on this, see “The Secrets of Great Teamwork,” HBR, June 2016.) At the same time, it has brought multiteaming—and the associated risks—to a whole new level. More and more people have at-will contracts and work not only on multiple projects but for multiple organizations. In many cases, companies are sharing team members’ time and smarts with market rivals.

Although most managers recognize the increasing prevalence of multiteaming, few have a complete understanding of how it affects their organizations, their teams, and individual employees. For instance, top leaders in one professional services

firm were surprised to learn who in their organization was most squeezed by multiteaming. First-year associates worked on as many as six projects in a week, which at a glance seemed like a lot. But the number rose steeply with tenure—employees worked on as many as 15 projects a week once they had reached the six-year mark. More-experienced people were members of fewer concurrent teams, but the more senior they got, the more likely they were to lead many projects at the same time. (See the exhibit “Who’s feeling the pain?”) Interviews revealed that working on multiple teams was stressful—one person likened it to being “slapped about” by different project leaders—despite benefits such as bringing lessons from one project to bear on others.

Who’s feeling the pain?

At one professional services firm, the employees most squeezed by multiteaming were mid-tenure associates—they helped with more and more projects as they gained experience. But the more senior people became, the more likely they were to lead many projects at the same time.

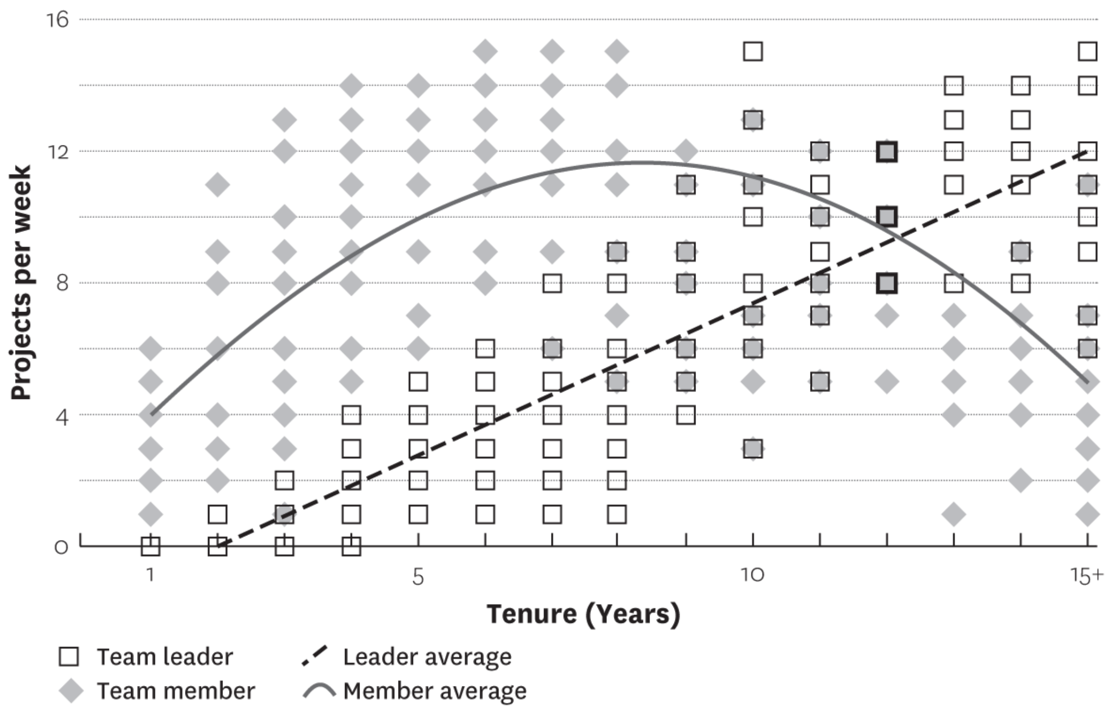

It's a classic "blind men and elephant problem." Managers see some of the benefits and some of the drawbacks firsthand but rarely all at once, because those things play out through different mechanisms and at different levels. Imagine, for example, a sales manager who wants to provide better solutions for customers by incorporating insights from her team members' experiences on other projects. That's not going to happen if splitting each individual's time across five projects means her team doesn't have the bandwidth to sit down and share those great ideas in the first place. Or consider a project manager who is thinking about adding a third engineer to his team—just 10% of a full-time equivalent—to reduce the load on his two overworked lead engineers. He may not recognize that this sort of slicing and dicing is the reason his first two engineers are in danger of

burnout—they are being pulled into too many competing projects. Examples like these abound.

For the most part, the benefits of multiteaming involve efficiency and knowledge flow, while the costs are largely intra- or interpersonal and psychological. That may be why the costs are tracked and managed less closely, if at all—and why they so often undermine the benefits without leaders' realizing it.

Managing the Challenges

Through our research and consulting, we have identified several ways that both team and organizational leaders can reduce the costs of multiteaming and better capitalize on its benefits. We'll outline them below.

Goals of multiteaming

(And the challenges that can undermine them)

Goals for teams	Challenges
Cost savings, because team members whose expertise is not required at the moment can bill their downtime to other projects	Weakened relationships and coherence within teams and projects Stress and burnout, particularly when members end up with assignments that exceed 100% time commitment
Process improvements as a result of importing	Interteam coordination costs so that schedules of projects with shared members don't collide

best practices and insights through shared members	Rocky transitions as members switch between *tasks* where their contributions are defined relative to other members' skills, adjust to different *roles* (boss on one team but subordinate on another), and learn new team *contexts* with unfamiliar routines, symbols, jokes, expectations, tolerance for ambiguity, and so on Reduced learning, because members lack time together to share knowledge and ideas Reduced motivation, because members have a small percentage of their time dedicated to any given project
Goals for organizations	**Challenges**
The capability to solve complex problems with members who have deep, specialized knowledge **Improved resource utilization** across projects (no one is dedicated to a project that needs only 5% of his or her time) **Increased knowledge transfer** and learning through shared membership	Politics and tensions over shared human resources Coordination costs of aligning timelines of projects even when they are not linked by content or workflow Weakened identification with the organization if people feel commoditized Increased risk as shocks affecting one team may pull shared members off other projects

Priorities for team leaders

Coordinating members' efforts (both within and across teams) and promoting engagement and adaptability are the key challenges for team leaders. Focusing on those goals early on, before your team even meets for the first time, will help you establish stronger relationships, reduce coordination costs, ease the friction of transitions, ward off political skirmishes, and identify risks so that you can better mitigate them. Here's how to do it:

Launch the team well to establish trust and familiarity. When fully dedicated to one team, people learn about their teammates' outside lives—family, hobbies, life events, and the like. This enables them to coordinate better (they know, for example, that one teammate is off-line during kids' bedtimes or that another routinely hits the gym during lunch). More important, it forges strong bonds and interpersonal trust, which team members need in order to seek and offer constructive feedback, introduce one another to valuable network connections, and rely on one another's technical expertise.

When multiteaming, in contrast, people tend to be hyperfocused on efficiency and are less inclined to share personal information. If you don't engineer personal interactions *for* them, chances are they'll be left with an anemic picture of their teammates, which can breed suspicion about why others fail to respond promptly, how committed they are to team outcomes, and so on. So make sure team members spend some time in the beginning getting to know their colleagues. This will

also help far-flung contributors give one another the benefit of the doubt later on. A Boston-based designer told us about his British counterpart:

> "I used to think that Sylvia was frosty and elitist, because she never jumped into our brainstorming sessions. Instead, she sent missives afterward, sometimes only to the project director. Then we spent a few days working together in person while I was in London, and I came to appreciate that she's an introvert who just needs time to process ideas before responding. Plus, because she had never met any of us, it was really hard for her to keep track of who had said what on the calls; she recognized only the leader's unique accent."

After the designer shared that "aha" with the team leader, the group switched to video calls so that everyone could see Sylvia's "thinking face" and she could feel confident that she was responding to the right people when making comments.

Formally launching the team—in person, if at all possible—helps a lot, especially if members open up about their own development goals. At McKinsey each team member, including the leader, explains how he or she expects to use that project to build or improve a critical skill. This level of openness not only encourages people to display some vulnerability (which is practically the definition of trust) but also gives members concrete ideas about how they can help one another.

The launch may feel like an unnecessary step if people know one another and everyone is ready to dive in, but research shows that team kickoffs can improve performance by up to 30%, in part because they increase peer-to-peer accountability. By clarifying roles and objectives up front and establishing group norms, you're letting people know what to expect from their colleagues. That's needed on any team, of course, but it's especially critical in organizations where people belong to several teams at once and must absorb *many* sets of roles, objectives, and norms to do good work across the board.

On teams that people frequently join or leave, you'll need to periodically "re-kick" to onboard new members and assess whether agreed-upon processes and expectations still make sense. A good rule of thumb is to do this whenever 15% of the team has changed.

Map everyone's skills. Figure out the full portfolio of capabilities that each person brings to the project—both technical skills and broader kinds of knowledge, such as familiarity with the customer's decision-making process, or a knack for negotiation, or insights about an important target market. Make sure everyone knows how each teammate contributes. This increases the chances that members will learn from one another. The pride people take in sharing their knowledge and the cohesion fostered by peer mentoring are often as valuable as the actual knowledge shared.

As with launching, it’s tempting to skip mapping if many members have worked together before. But we’ve found that even familiar teams are likely to hold outdated assumptions about individuals’ potential contributions and often disagree about their teammates’ expertise. As a result, they may argue about which roles members should play or bristle at assignments, thinking they’re unfair or a bad fit. People may also waste time seeking outside resources when a teammate already has the needed knowledge, which demotivates those whose skills have been overlooked.

Sherif, a tax expert, experienced these problems when he joined with four colleagues to pitch a new client. “We’d all worked together on prior projects over the years—enough, we assumed, to know one another’s ‘sweet spots,’” he told us. “Over time, though, I grew more and more frustrated that two of my partners kept adding bits of regulatory advice to the pitch document—that’s why I was on the team! I was handling nearly the exact same issue for a current client. I felt undermined, and the more they tried to sideline me, the more cantankerous I got.” A few days before the client meeting, the group talked it out and discovered that Sherif had been honing his specialist expertise on projects the others hadn’t been part of. They simply didn’t realize what he had to offer. “We’d all been running in so many directions at the same time that our individual knowledge was changing quickly,” he says. “No wonder we had friction.”

Skills mapping could have prevented this. It also streamlines communication (no need to "reply all" if you know who's actually responsible for an issue). And it equips members to hold one another accountable for high-quality, on-time delivery, which is otherwise tricky when people are frequently coming and going. Creating the expectation of peer accountability relieves you as the team leader from some of that day-to-day oversight, freeing you up to scan the environment for potential shocks from other teams, for example, or to handle some of the inevitable negotiations about shared resources.

Manage time across teams. As you form a team, explicitly talk about everyone's competing priorities up front. By preemptively identifying crunch periods across projects, you can revamp deadlines or plan on spending more hands-on time yourself at certain points. Making the topic "discussable" so that people won't feel guilty about conflicts allows the team to openly and productively handle these issues when they come up later.

Establishing the right rhythm of meetings will make it easier to manage time across teams and address competing priorities. At the outset, you'll want to schedule several full-team meetings at critical junctures. (Research shows, for instance, that the halfway point in any project is a vital moment for a check-in, because that's when people shift into a higher gear, acutely aware that their time is limited.) Make attendance truly mandatory, and ensure it by giving each team member a piece of the meetings to run—even if it's just for 10 minutes. Check in early to see that all

members have cleared meeting dates with their other teams. Ideally, the organizational culture will support formal check-in meetings as a high priority. If not, you may need to coordinate with other team leaders before putting a schedule together.

When you plan other team meetings, invite exactly who's needed and no one else, to minimize scheduling conflicts with other teams. Most of the time, you won't need everyone. Meet in subteams whenever possible. Don't forget to leverage technology: Instead of using precious live meeting time for updates, send a three-line e-mail or keep an online dashboard updated so that people can track progress as needed. Although technology doesn't replace face-to-face interaction, it can tide you over when a full meeting is too costly. And be creative: Younger team members are more likely to watch a 30-second video update than to read a two-page memo. Brief, spontaneous check-ins with team members over Skype or FaceTime can keep you updated on their competing deadlines; this visual interaction makes it more likely that you'll pick up cues about their stress and motivation levels, too.

Create a learning environment. Learning makes work feel more meaningful, and it's supposed to be a major benefit of multiteaming—but it often gets crowded out by time pressures. There are other obstacles as well: Even if you've worked to build trust and personal connections, it's harder for multiteamers to give effective feedback than it is for dedicated team members, because people whose time is divided among several projects are

less likely to regularly observe their teammates' actions or to be present at a time that "feels right" to offer critiques. Members who see only a small slice of a project may lack the context to fully understand what kind of feedback is appropriate. They also tend to focus on short-term tasks and to communicate with one another only when required.

Carrie, for example, was promoted to run the development office of a major metropolitan hospital, and her new 20-person staff was splitting its time among dozens of projects each week. After six months she realized, "We were all living in a feedback desert. I literally hadn't had a single comment in half a year about how I could do my job better, despite clear examples of projects that hadn't lived up to expectations." To change the tone, she modeled seeking input and responding to it constructively. "Doing so day in and day out, I started to create an environment where people shared their concerns to get help as soon as they needed it," she says. "Over time, it felt safe enough to put in more-formal processes to review projects and allow everyone to learn from errors without fear of retribution or blame."

You can also designate team members from different functions or offices to colead parts of the project so that they benefit from greater cross-contact; a formal assignment makes it more likely that they'll devote time to learning from each other. Similarly, pair a highly experienced team member with someone more junior and help them understand what both can gain from the

exchange—it's not just one-way learning flowing down to the junior person.

Foster curiosity by posing "What if . . . ?" questions when it's likely that different members' backgrounds will provide new insights. If you get a question that you know another member could answer more fully, given his or her experience, redirect the asker and prompt the expert to do a bit of tutoring.

Boost motivation. On traditional, fixed teams, a strong sense of cohesion and group identity motivates members. But leaders in multiteaming environments need to leverage more of an exchange relationship. The ability to get jazzed about a project naturally flags when members spend only a small amount of time on it. Their inner accountant asks, "If I'll get only 10% of the credit, how much time and effort should I devote to this?" Figure out what your ten-percenters really value and frame the work in terms of those rewards. For example, if you have a Millennial who is eager to develop transferable skills, you might occasionally take time during meetings to have team members share and learn something new, or hold a workshop at the end of the project in which members cross-train.

Remember, too, that a sense of fairness drives many behaviors. If people feel they are pulling their weight while others slack off, they quickly become demotivated. When team members are tugged in many directions, it's often difficult for each one to recognize and appreciate how hard the others are working. As the leader, keep publicly acknowledging various members'

contributions so that they become visible to the whole team, spawning a greater awareness of the collective efforts.

Like Christine, the frustrated leader of the Analytix software team, you might be feeling the strain of sharing valuable talent with other teams. Before you reach the breaking point, take these steps to clarify and manage your interdependency with other teams. They will help you avoid conflicts when that's possible, defuse them when it's not, and set an example of better collaboration with other team leaders—peers who face the same challenges you do.

Priorities for organizational leaders

If you're leading an organization where multiteaming is prevalent, you'll need to keep a close eye on how—and how many—members are shared across teams. We've found that you can reduce organizational risk and boost innovation by following these steps:

Map and analyze human capital interdependence. Patterns of team overlap range from highly concentrated (a large proportion of members are shared by just a few teams) to highly dispersed (the sharing is spread out across many teams).

Each pattern has its own implications for risk management. When a surprise problem jolts one team, the cry "All hands on deck" pulls shared members off their other teams—with disproportionately large effects on teams that have a concentrated overlap in members. When the overlap is more

dispersed, the shock will be felt by more teams but to a lesser extent by each one. (See the exhibit “Who takes the hit?”)

Who takes the hit?

When a couple of teams share many members, a shock to one group severely jolts the other, because people shift their efforts from ongoing work to firefighting.

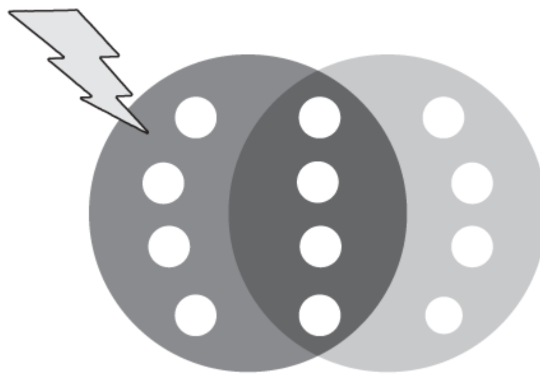

When many teams share just one or two members, a shock to one group has a minor impact on the others—but the effects ripple throughout the organization.

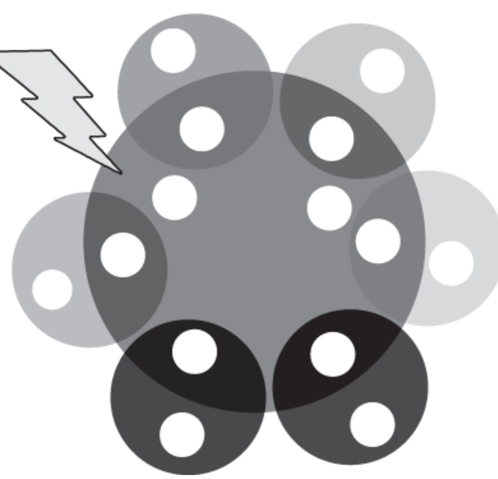

There are implications for knowledge transfer as well. Best practices travel from one project to the next as team members share what's working—and what isn't—on their other projects. Highly concentrated overlap makes it easier to spread ideas from one team to another; highly dispersed overlap makes it easier to spread them to more teams.

Keep an accurate map of the links among teams in your organization through periodic updates from managers and team members. The frequency of these check-ins will depend on the life cycles of your teams. You'll need them more often if teams and assignments change week to week, less often if you've got yearlong projects with stable membership. This bird's-eye view will help you see which teams fail to pick up on new trends because they're too isolated, for instance, and which are so tightly interconnected that they aren't mitigating the risks of their shared membership.

The question we get most often about mapping interdependence is "What's the right amount?" Unfortunately, there's no magic answer—either for overlap between teams or for the number of teams per individual. Both targets depend highly on context. When teams are very similar in their tasks and culture, transitioning between them is relatively easy, so you can have a large amount of overlap and members can be on more of them. Transitioning across teams with very different tasks or cultures should be kept to a minimum, however—it's a bigger, costlier shift. Interestingly, the reverse holds true when

workloads differ across teams, because members aren't in high demand from all teams at the same time (they aren't as susceptible to burnout as, say, tax advisers in April are).

Once you've done all this analysis, it's time to address the shortcomings you've uncovered—which brings us to the next two steps.

Promote knowledge flows. Pay close attention to teams that share few or no members with others—whether that's by design or by accident. These "islands" will require help staying informed about what's working elsewhere in the organization, sharing their knowledge and ideas, and deciding who would be the best resource to apply to a given task.

Your goal here is to establish knowledge transfer as a cultural norm, which involves getting employees to recognize that everyone wins when they take the time to share insights across projects. As with any cultural shift, it's important to lead by example and to reward those who follow suit. That's simple to say—but not so simple to do. To make it easier, highlight the benefits of sharing, and provide processes and technology to facilitate it, such as brown-bag lunches and online forums. One tech firm we worked with made a point of celebrating project breakthroughs that were attributed to transferred best practices. R&D teams at a manufacturing company shared monthly testimonials from individuals who had gained new insights through cross-staffing. In both cases the objective was to make the benefits of knowledge transfer clear—and to counter the

ever-present pressure for people to keep their heads down and focus on immediate tasks.

Buffer against shocks. How can you prevent shocks in one team from being transmitted to others? Often you can't—but knowing how teams are connected through shared membership allows you to anticipate *where* some shocks may be transferred and to design small amounts of slack into the system to absorb them. This doesn't mean having people sit around twiddling their thumbs just in case. Rather, you're enabling them to shift their attention when needed. One engineering firm we worked with had identified several skilled "firefighters" and assigned them to long-term projects that wouldn't suffer if they had to address urgent problems elsewhere. This had the added benefit of providing those individuals with exciting challenges that were a welcome change of pace from their day-to-day work.

It takes a critical eye and a clear set of strategic priorities to determine which projects can be disrupted and which can't. Sometimes it makes sense to give certain projects "protected" status, exempting members of those teams from answering others' firefighting calls. Overall, the idea is to be responsive to immediate problems without sacrificing teams' ongoing needs. Of course, even if you've built slack into team design, you may occasionally have to jump in with extra resources to save critical projects that take a hit. But your other teams will feel less pain when you do.

None of this is easy. You may need to work with HR or IT to establish processes or systems that will allow you to track multiteaming more accurately across the organization. You may even need to create a new role to define and coordinate these efforts effectively. And people may resist the increased oversight—it can feel like micromanagement to team leaders and members who are accustomed to having freer rein, particularly in entrepreneurial cultures. Still, in the end such investments are worthwhile; it's actually more costly to allow the trade-offs of multiteaming to go unchecked. If you're open about the problems you're trying to solve with all this transparency, people are less likely to feel surveilled or constrained by it and more likely to see the upside.

Nearly every knowledge worker these days is a member of multiple concurrent teams. Together, organizational and team leaders can make the most of that trend by creating an environment where multiteamers will thrive. Some of this involves managing interdependence risks, articulating and navigating groups' competing priorities, and removing obstacles to strategic coordination across groups. And some entails building stronger connections and greater trust among people who spend only a small fraction of their time together.

All around, it's a significant investment of time and effort. But organizations pay a much higher price when they neglect the costs of multiteaming in hot pursuit of its benefits.

Originally published in September–October 2017. Reprint R1705C

Why Do We Undervalue Competent Management?

by Raffaella Sadun, Nicholas Bloom, and John Van Reenen

IN MBA PROGRAMS, students are taught that companies can't expect to compete on the basis of internal managerial competencies because they're just too easy to copy. Operational effectiveness—doing the same thing as other companies but doing it exceptionally well—is not a path to sustainable advantage in the competitive universe. To stay ahead, the thinking goes, a company must stake out a distinctive strategic position—doing something different than its rivals. This is what the C-suite should focus on, leaving middle and lower-level managers to handle the nuts and bolts of managing the organization and executing plans.

Michael Porter articulated the difference between strategy and operational effectiveness in his seminal 1996 HBR article, "What Is Strategy?" The article's analysis of strategy and the strategist's role is rightly influential, but our research shows that simple

managerial competence is more important—and less imitable—than Porter argued.

If you look at the data, it becomes clear that core management practices can't be taken for granted. There are vast differences in how well companies execute basic tasks like setting targets and grooming talent, and those differences matter: Firms with strong managerial processes perform significantly better on high-level metrics such as productivity, profitability, growth, and longevity. In addition, the differences in the quality of those processes—and in performance—persist over time, suggesting that competent management is not easy to replicate.

Nobody has ever argued that operational excellence doesn't matter. But we contend that it should be treated as a crucial complement to strategy—and that this is true now more than ever. After all, if a firm can't get the operational basics right, it doesn't matter how brilliant its strategy is. On the other hand, if firms have sound fundamental management practices, they can build on them, developing more-sophisticated capabilities—such as data analytics, evidence-based decision making, and cross-functional communication—that are essential to success in uncertain, volatile industries.

Achieving managerial competence takes effort, though: It requires sizable investments in people and processes throughout good times and bad. These investments, we argue, represent a major barrier to imitation.

In this article we'll review our research findings and then discuss the obstacles that often prevent executives from devoting sufficient resources to improving management skills and practices. Throughout, we'll show that such investments are a powerful way to become more competitive. If the world has really entered a "new normal" of low productivity growth, as Robert Gordon and others have argued, pushing managerial capital up a level could be the best route out of the performance doldrums.

The Research

Over the past century, scholars have learned a great deal about how core management processes affect a company's performance. For example, researchers such as Kim Clark, Bob Hayes, and David Garvin documented differences within factories, industries, and companies. But a lack of big data encompassing many firms, industries, and countries inhibited the statistical study of management practices. In the past decade, however, we have developed ways to robustly measure core management practices, and we can now show that their adoption accounts for a large fraction of performance differences across firms and countries.

Idea in Brief

The Conventional Wisdom

It's a truism among strategists that you can't compete on the basis of better management processes because they're easily copied. Operational excellence is table stakes in the competitive marketplace.

What the Data Shows

There are three problems with this thinking. First, effective management processes are highly correlated with measures of strategic success. Second, differences in process quality persist over time. Third, there's little evidence that best-in-class processes can be imitated. GM tried for years to adopt Toyota's superior production system and failed miserably.

Implications

Organizations need competent management just as much as they need analytical brilliance. We should stop teaching business school students that operational issues are beneath the CEO—and should encourage firms to invest in strengthening management throughout the organization.

As we've described in earlier articles in HBR, in 2002 we began an in-depth study of how organizations in 34 countries use (or don't use) core management practices. Building on a survey instrument that was initially developed by John Dowdy and Stephen Dorgan at McKinsey, we set out to rate companies on their use of 18 practices in four areas: operations management, performance monitoring, target setting, and talent management. (See the sidebar "Core Managerial Practices" for a detailed list. Though these don't represent the full set of important managerial practices, we have found that they're good proxies for general operational excellence.) The ratings ranged from poor to nonexistent at the low end (say, for performance monitoring

using metrics that did not indicate directly whether overall business objectives were being met) to very sophisticated at the high end (for performance monitoring that continuously tracked and communicated metrics, both formally and informally, to all staff with an array of visual tools).

Core Managerial Practices

In our research, we assess the sophistication with which organizations manage the four broad dimensions—and the 18 specific aspects—of management shown below. The list varies slightly depending on sector (this one is for manufacturers). It's not exhaustive, but companies that manage these fundamentals well tend to have high levels of overall operational excellence.

Operations Management

- Use of lean techniques
- Reasons for adopting lean processes

Performance Monitoring

- Process documentation
- Use of key performance indicators
- KPI reviews
- Discussion of results
- Consequences for missing targets

Target Setting

- Choice of targets

- Connection to strategy, extent to which targets cascade down to individual workers
- Time horizon
- Level of challenge
- Clarity of goals and measurement

Talent Management

- Talent mindset at the highest levels
- Stretch goals
- Management of low performance
- Talent development
- Employee value proposition
- Talent retention

Our aim was to gather reliable data that was fully comparable across firms and covered a large, representative sample of enterprises around the world. We realized that to do that, we needed to manage the data collection ourselves, which we did with the help of a large team of people from the Centre for Economic Performance at the London School of Economics. To date the team has interviewed managers from more than 12,000 companies about their practices. On the basis of the information gathered, we rate every organization on each management practice, using a 1 to 5 scale in which higher scores indicate greater adoption. Those ratings are then averaged to produce an

overall management score for each company. (For more details, see the sidebar “About the Research.”)

About the Research

Our research project, World Management Survey, has examined the adoption and use of management practices across more than 12,000 firms and 34 countries. We measure each organization’s performance on 18 specific practices in four areas: operations management, performance monitoring, target setting, and talent management. To do that, we have experienced interviewers speak by phone with a firm’s plant managers, asking everyone the same 18 open-ended questions and following up with more questions until they have a good sense of the firm’s habits. A listener, who doesn’t have information about the organization’s financial performance, independently scores the organization on each question and each practice.

So far we’ve conducted more than 20,000 interviews and surveyed companies in four sectors: manufacturing, health care, retail, and higher education. More information about our methodology is available on our website, worldmanagementsurvey.com, where readers can also download the survey, fill in their own responses, and compare their organizations against the benchmarks in our data set. Obviously, the results won’t be as complete, or as trustworthy, as they’d be if the organization were being independently assessed, but the process can provide a useful broad-strokes view.

That data has led us to two main findings: First, achieving operational excellence is still a massive challenge for many organizations. Even well-informed and well-structured companies often struggle with it. This is true across countries

and industries—and in spite of the fact that many of the managerial processes we studied are well known.

The dispersion of management scores across firms was wide. Big differences across countries were evident, but a major fraction of the variation (approximately 60%) was actually within countries. (See the exhibit "Management quality varies across—and within—countries.") The discrepancies were substantial even within rich countries like the United States.

Management quality varies across—and within—countries

Some countries get higher average ratings than others on the use of management processes. But as data from this sample of countries shows, in-country variation is even more striking. The black bars indicate what percentage of firms in each country fell into each scoring range (1 equaled the worst and 5 the best performance). The gray bars show total global percentages.

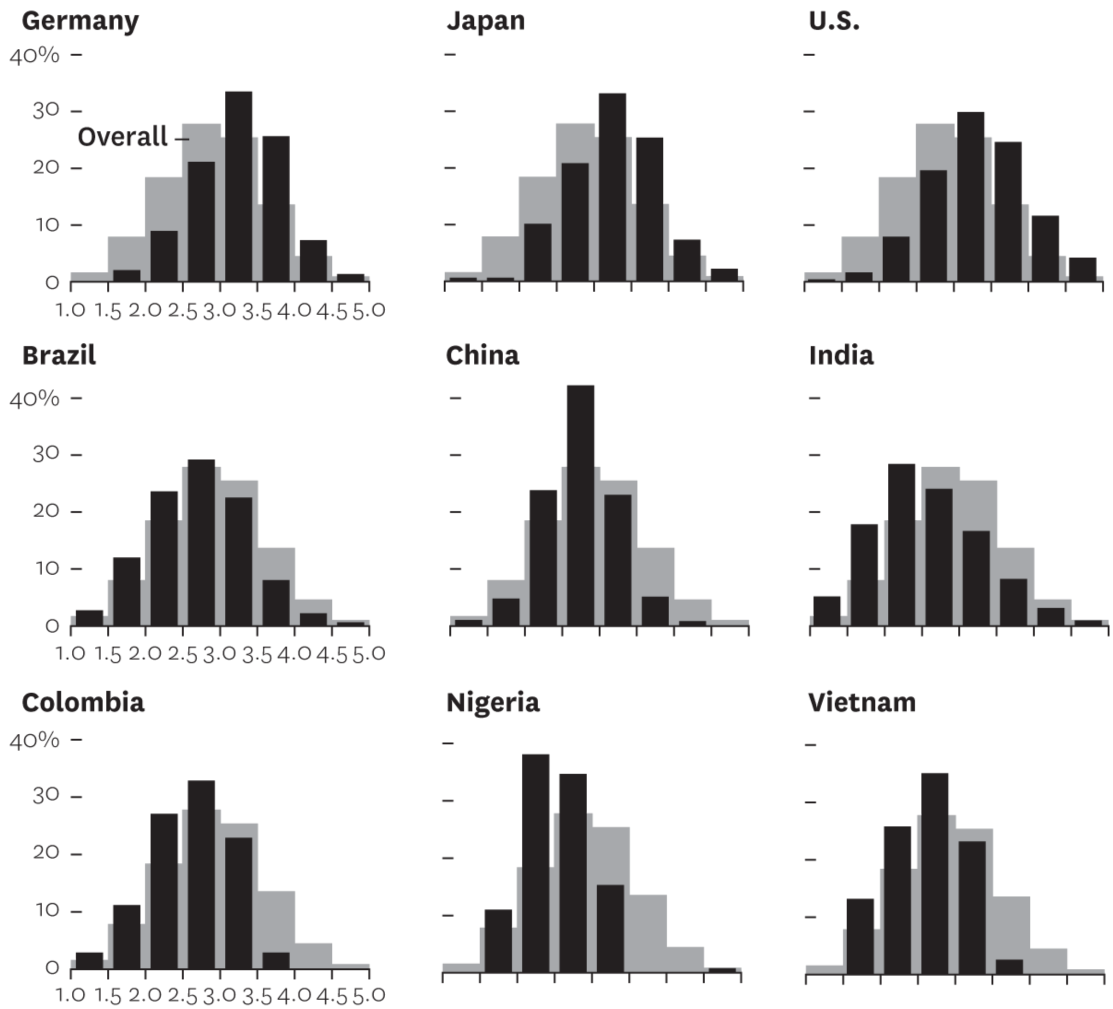

In our entire sample we found that 11% of firms had an average score of 2 or less, which corresponds to very weak monitoring, little effort to identify and fix problems within the organization, almost no targets for employees, and promotions and rewards based on tenure or family connections. At the other end of the spectrum we identified clear management superstars across all the countries surveyed: Six percent of the firms in our sample had an average score of 4 or greater. In other words they had rigorous performance monitoring, systems geared to optimize

the flow of information across and within functions, continuous improvement programs that supported short- and long-term targets, and performance systems that rewarded and advanced great employees and helped underperformers turn around or move on.

By interviewing several companies multiple times throughout the past decade, we were able to observe that these large differences in the adoption of core management practices were long-lasting. This isn't really surprising: According to our estimates, the costs involved in improving management practices are as high as those associated with capital investments such as buildings and equipment.

One of our findings may surprise readers: These differences show up within companies, too. A project conducted with the U.S. Census revealed that variations in management practices inside firms across their plants accounted for about one-third of total variations across all plant locations. This was particularly true in large firms, where practices can differ a great deal across plants, divisions, and regions. Even the biggest and most successful firms typically fail to implement best practices throughout the whole organization. Some parts of it are effectively managed, but other parts struggle.

Our second major finding was that the large, persistent gaps in basic managerial practices we documented were associated with large, persistent differences in firm performance. As we've noted, our data shows that better-managed firms are more profitable,

grow faster, and are less likely to die. Indeed, moving a firm from the worst 10% to the best 10% of management practices is associated with a $15 million increase in profits, 25% faster annual growth, and 75% higher productivity. Better-managed firms also spend 10 times as much on R&D and increase their patenting by a factor of 10 as well—which suggests that they're not sacrificing innovation to efficiency. They also attract more talented employees and foster better worker well-being. These patterns were evident in all countries and industries. (For a sample of metrics, see the exhibit "Good management correlates with strong performance.")

Good management correlates with strong performance

The companies scoring in the top decile on management outperformed on a variety of strategic measures. Performance by decile:

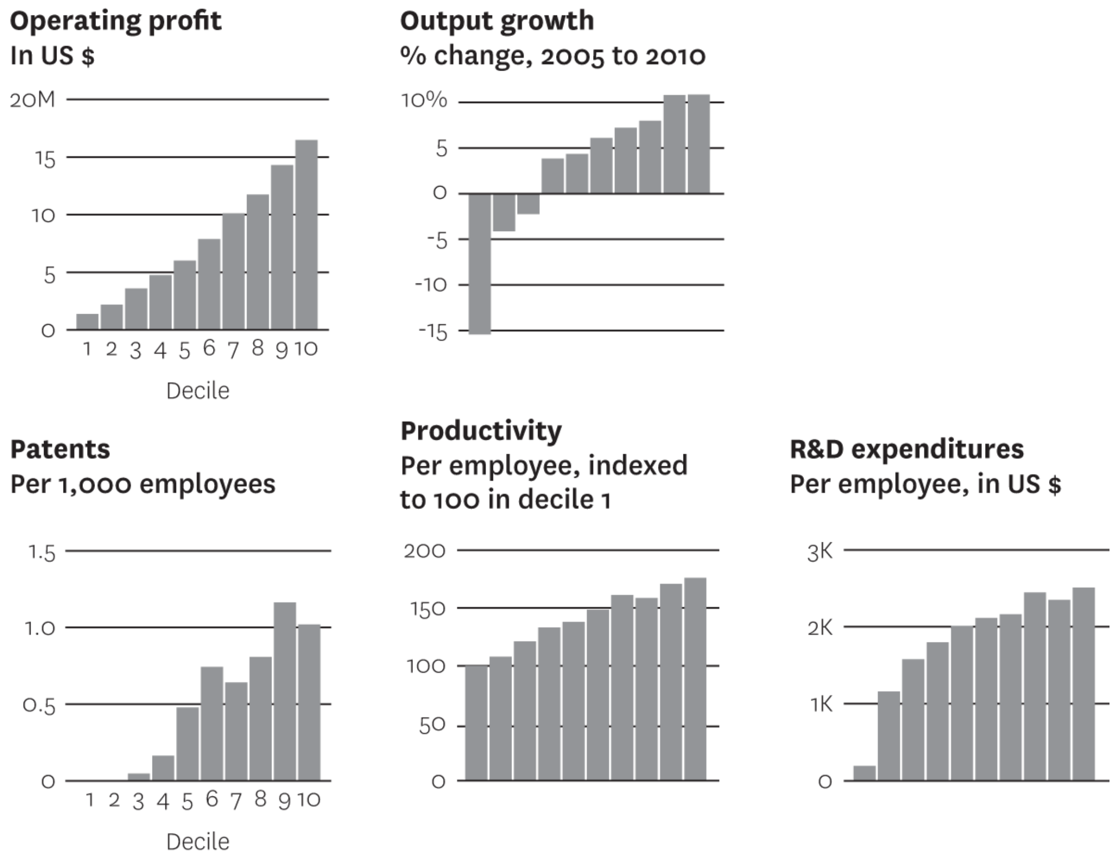

But these empirical findings raise a major question: If the benefits of core managerial practices are really so large and extensive, why doesn't every company focus on strengthening them? Also, a more existential issue (which we'll address toward the end of the article) is, What should executives, business schools, and policy makers take away from this body of research?

What Causes the Differences?

Some of the variation in management practice is driven by external factors. The intensity of competition is one; competition creates a strong incentive to reduce inefficiencies and kills off

badly managed firms. Labor regulations play a role as well; they can make it difficult to give opportunities to employees on the basis of merit or to adopt performance-related compensation. On the flip side, regulators may be in a position to create incentives for employee training or support firms that prioritize managerial competence.

We've also observed that inconsistencies often result from stubborn blind spots and deficiencies within companies. Here are the things that typically hinder the adoption of essential management practices:

False perceptions

Our research indicates that a surprisingly large number of managers are unable to objectively judge how badly (or well) their firms are run. (Similar biases show up in other settings. For example, 70% of students, 80% of drivers, and 90% of university teachers rate themselves as "above average.")

Consider the average response we got to the question "On a scale from 1 to 10, how well managed is your firm?," which we posed to each manager at the end of the survey interview. (See the exhibit "Overconfidence is a problem for managers.") Most managers have a very optimistic assessment of the quality of their companies' practices. Indeed, the median answer was a 7. Furthermore, we found zero correlation between perceived management quality and actual quality (as indicated by both their firms' management scores and their firms' performance), suggesting that self-assessments are a long way from reality.

Overconfidence is a problem for managers

At the end of every interview, we ask managers to say how well they think their organizations are run and to score them on a scale from 1 (worst) to 10 (best). Overall, their responses are far more positive than warranted.

This large gap is problematic, because it implies that even managers who really need to improve their practices often don't take the initiative, in the false belief that they're doing just fine.

In a variant of this problem, managers may overestimate the costs of introducing new practices or underestimate how much difference they could make. This was a situation we encountered in a field experiment that one of us conducted with 28 Indian textile manufacturers. Accenture had been hired by a Stanford–

World Bank project to improve their management practices, but many proposed enhancements—such as quality control systems, employee rewards, and production planning—were not implemented because of skepticism about their benefits. Consultants trying to introduce methods that are standard in most U.S. or Japanese factories were met with claims that “it will never work here” or “we do things our way.” Yet the firms that adopted the methods boosted their performance.

Perception problems are hard but not impossible to eradicate. The key is to improve the quality of information available to managers so that they have an objective way to evaluate their relative performance.

As our survey shows, self-reported metrics are likely to be at best very noisy—they’re imperfect indicators of what really happens on the ground. There are various reasons why. A common issue is that employees don’t raise problems for fear of being blamed for those they identify. That dynamic deprives managers of critical knowledge needed to understand a firm’s gaps.

In our experience, managers can address this issue by proactively creating opportunities for candid—and blame-free—discussions with their employees. That’s the approach followed by Danaher, a large U.S. conglomerate known for its relentless (and effective) adoption of the Danaher Business System (DBS)—a tool kit of managerial processes modeled on the Toyota Production System—across its many subsidiaries. Danaher

typically initiates the relationship with a newly acquired subsidiary through a series of hands-on, structured interactions between senior Danaher managers and the acquisition's top executives, which challenge the latter to identify managerial gaps that may be preventing the business from fulfilling its potential. People taking part in these open conversations—especially those with longer tenure—describe them as eye-opening experiences that significantly change attitudes toward core management processes.

Governance structure

In other cases, managers may be fully aware of the need to improve their practices but pass on this opportunity for fear that change may jeopardize private objectives. This problem is particularly common in firms that are owned and run by families, as you can see in the exhibit "Family-run firms tend to have weaker management." Even when we cut the data by firm size, sector of activity, and country, family-run enterprises still had the lowest average management scores.

Family-run firms tend to have weaker management

Average score by type of ownership (1 = worst; 5 = best)

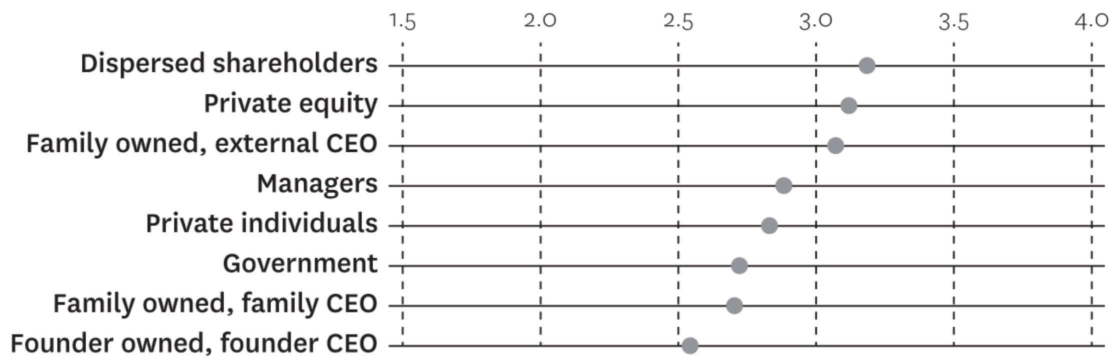

Why are family firms so reluctant to embrace strong management processes? One explanation—which finds support in our research—is that their adoption may have significant personal costs to family members. New practices may require hiring or delegating authority to talent outside the family circle. (Indeed, we've seen that higher management scores tend to go hand-in-hand with more-decentralized decision making.)

An example of this is Gokaldas Exports, a family- owned business founded in 1979 that had grown into India's largest apparel exporter by 2004. Gokaldas was a highly successful firm with 30,000 workers, was valued at approximately $215 million, and exported nearly 90% of its production. Its founder, Jhamandas Hinduja, had bequeathed control of the company to three sons, each of whom brought his own son into the business. Nike, a major customer, wanted Gokaldas to introduce lean management practices; it put the company in touch with consultants who could help make that happen. Yet the CEO was resistant. It took rising competition from Bangladesh, multiple visits to see lean manufacturing in action at firms across Asia and

the United States, and finally the intervention of other family members (one of whom we taught in business school) to overcome his reluctance.

Self-reflection exercises can help family CEOs clarify whether they value their firms' long-term success more than "being the boss"—even if success means sharing the glory with other managers. In our experience a candid evaluation of one's priorities is crucial—managers are often oblivious to the fact that their own desire for control may be inhibiting the growth and success of their organizations.

In addition, family executives—and especially owners—should understand that introducing new managerial capabilities within the firm does not necessarily entail a loss of control. It is more likely to create a different role for them—but not necessarily fewer responsibilities.

That is what happened at Moleskine, based in Milan, Italy. Launched in 1997 by three friends, Moleskine went from being a niche notebook producer to a market leader in the space of a few years. Its success created a dilemma for its founders: While it was clear that the company had tremendous potential to grow further, they also recognized the pressing need to professionalize its operations. The founders searched for a private equity firm that could provide the necessary capital and expertise and help them find a new CEO. Eventually, they chose Syntegra Capital and Arrigo Berni, an experienced chief executive who had held leadership roles at family-owned producers of luxury products.

Berni brought new rigor to strategy development and operations and at the same time crafted a role for the founders that made the most of their commercial and design expertise. Thanks to this successful partnership—and an IPO in 2013—Moleskine was able to deepen its competitive advantage and develop new growth opportunities globally.

Skill deficits

Good management practices require capabilities (such as numeracy and analytical skills) that may be lacking in a firm's workforce, especially in emerging economies. Indeed, our data shows that the average management score is significantly higher at firms with better-educated employees. Being located near a leading university or business school is also strongly associated with better management scores. Superior performance is likelier when executive education can be had nearby, it seems. While to some extent the availability of skills is shaped by a firm's specific context, managers can play a critical role by recognizing the importance of employees' basic skills and providing internal training programs.

Organizational politics and culture

Even when top managers correctly perceive what needs to be done, are motivated to make changes, and have the right skills, the adoption of core management processes can be a challenge. Videojet, a subsidiary acquired by Danaher, provides a case in point. In 2005, Videojet launched a new internal initiative that

required the engineering and sales teams to collaborate on developing an innovative printer. The Videojet executives decided to use core DBS managerial processes—which up to that point had been used almost exclusively within manufacturing—to structure regular debriefing and problem-solving sessions between the two teams.

Unfortunately, preexisting divides between engineers and salespeople meant that the structured interactions, which had been effective in driving continuous improvement in manufacturing, became perfunctory meetings. For example, just before the product launch, a member of the sales team raised concerns about some technical aspects of the new printer, which in his eyes could seriously compromise its success. The core DBS processes had been introduced to help teams identify and address precisely this type of concern. Whereas in manufacturing, employees were encouraged to stop the production line to flag quality problems in real time so that they could be isolated and fixed, in this instance the feedback was ignored and interpreted by the rest of the team as a boycotting attempt rather than a constructive suggestion. Shortly after this episode, the printer was launched to a poor market reception, which confirmed the gravity of the issues the salesperson had raised. Thanks to this experience, Videojet executives understood that they would need to work more consciously to foster interactions between diverse pockets of expertise within the firm. They continued to use the DBS tools but also

committed to frequent, longer structured interactions and collective sign-offs between engineers and salespeople during the various product development stages. Videojet launched a very successful printer just a couple of years after the initial failed product launch and has since become an exemplar in the use of DBS tools for product development.

Sometimes the organization at large resists change. Susan Helper and Rebecca Henderson provide a fascinating account of the difficulties GM encountered in implementing the Toyota Production System during the 1980s and 1990s. Even in the face of mounting competition, GM found it hard to adopt Toyota's superior management methods, mainly because of adversarial relationships with suppliers and blue-collar workers. Employees, for example, thought that any productivity enhancement from the new practices would just lead to head-count reductions and would more generally put employees under greater pressure. This distrust inhibited GM's ability to negotiate for the working arrangements needed to introduce the new practices (such as teams and joint problem solving).

Videojet's and GM's experiences illustrate a fundamental issue: Management practices often rely on a complicated shared understanding among people within the firm. The inability to foster it can easily kill the efforts of the most able and well-intentioned managers. On the other hand, once such an understanding is in place, it's very difficult for competitors to replicate.

A question that managers face is how to create this common understanding. Changing individual incentives is unlikely to work, since the adoption of new processes usually requires the cooperation of teams of people; it's difficult to disentangle the rewards to be assigned to a single employee. And adoption is hard to measure, so it would be challenging to tie an individual bonus to the implementation of a certain practice. As organizational economists know, simple contractual solutions are hardly effective in these situations.

But managers have a different weapon at their disposal, which in our experience can potentially be more effective. It's their presence. The successful adoption stories that we've encountered in our research often took place in organizations where someone very high up signaled the importance of change through personal involvement, constant communication, message reinforcement, and visibility. "Walking the talk" matters enormously and can drastically affect the odds of success for change initiatives.

This idea is supported by a large-scale research project on the relationship between management and CEO behavior that Raffaella conducted with a different team of researchers at the London School of Economics and Columbia University. After a painstaking exercise in which they codified the agendas of more than 1,200 CEOs of manufacturing firms in six countries, they found that management quality was significantly higher in

organizations in which CEOs dedicated a larger portion of their time to employees than to outside stakeholders.

Though core management practices may appear to be relatively simple—in that they often rely on nontechnological investments—they are not light switches that can be flipped on and off at will. They require a profound commitment from the top, an understanding of the types of skills required for adoption, and—ultimately—a fundamental shift in mentality at all levels of the organization.

Next Steps

Our findings have implications for how managers are trained. Today business students are encouraged to judge case studies about operational effectiveness as "nonstrategic" and to see these issues as not pertinent to the role of the CEO. But it's unwise to teach future leaders that strategic decision making and basic management processes are unrelated, and that the first is far more important to competitive success than the second.

Indeed, our work suggests that the management community may have badly underestimated the benefits of core management practices—as well as the investment needed to strengthen them—by relegating them to the domain of "easy to replicate." Managers should certainly dedicate their time to fundamental strategic choices, but they should not suppose that fostering strong managerial practices is below their pay grade. Just as the ability to discern competitive shifts is important to

firm performance, so too is the ability to make sure that operational effectiveness is truly part of the organization's DNA.

One frequent suggestion in this era of flattened organizations is that everyone has to be a strategist. But we'd suggest that everyone also needs to be a manager. Core management practices, established thoughtfully, can go a long way toward plugging the execution gap and ensuring that strategy gets the best possible chance to succeed.

Originally published in September–October 2017. Reprint R1705K

“Numbers Take Us Only So Far”

by Maxine Williams

I WAS ONCE EVICTED from an apartment because I was black. I had secured a lovely place on the banks of Lake Geneva through an agent and therefore hadn’t met the owner in person before signing the lease. Once my family and I moved in and the color of my skin was clear to see, the landlady asked us to leave. If she had known that I was black, I was told, she would never have rented to me.

Terrible as it felt at the time, her directness was useful to me. It meant I didn’t have to scour the facts looking for some other, nonracist rationale for her sudden rejection.

Many people have been denied housing, bank loans, jobs, promotions, and more because of their race. But they’re rarely told that’s the reason, as I was—particularly in the workplace. For one thing, such discrimination is illegal. For another, executives tend to think—and have a strong desire to believe—that they’re hiring and promoting people fairly when they aren’t. (Research shows that individuals who view themselves

as objective are often the ones who apply the most unconscious bias.) Though managers don't cite or (usually) even perceive race as a factor in their decisions, they use ambiguous assessment criteria to filter out people who aren't like them, research by Kellogg professor Lauren Rivera shows. People in marginalized racial and ethnic groups are deemed more often than whites to be "not the right cultural fit" or "not ready" for high-level roles; they're taken out of the running because their "communication style" is somehow off the mark. They're left only with lingering suspicions that their identity is the real issue, especially when decision makers' bias is masked by good intentions.

I work in the field of diversity. I've also been black my whole life. So I know that underrepresented people in the workplace yearn for two things: The first is to hear that they're not crazy to suspect, at times, that there's a connection between negative treatment and bias. The second is to be offered institutional support.

The first need has a clear path to fulfillment. When we encounter colleagues or friends who have been mistreated and who believe that their identity may be the reason, we should acknowledge that it's fair to be suspicious. There's no leap of faith here—numerous studies show how pervasive such bias still is.

But how can we address the second need? In an effort to find valid, scalable ways to counteract or reverse bias and promote

diversity, organizations are turning to people analytics—a relatively new field in business operations and talent management that replaces gut decisions with data-driven practices. People analytics aspires to be "evidence based." And for some HR issues—such as figuring out how many job interviews are needed to assess a candidate, or determining how employees' work commutes affect their job satisfaction—it is. Statistically significant findings have led to some big changes in organizations. Unfortunately, companies that try to apply analytics to the challenges of underrepresented groups at work often complain that the relevant data sets don't include enough people to produce reliable insights—the sample size, the *n,* is too small. Basically they're saying, "If only there were more of you, we could tell you why there are so few of you."

Companies have access to more data than they realize, however. To supplement a small *n,* they can venture out and look at the larger context in which they operate. But data volume alone won't give leaders the insight they need to increase diversity in their organizations. They must also take a closer look at the individuals from underrepresented groups who work for them—those who barely register on the analytics radar.

Idea in Brief

Though executives tend to think—and want to believe—they're hiring and promoting fairly, bias still creeps into their decisions. They often

use ambiguous criteria to filter out people who aren't like them or deem people from minority groups to be "not the right cultural fit," leaving those employees with the uneasy feeling that their identity might be the real issue.

Companies need to acknowledge that it's fair for employees from underrepresented groups to be suspicious about bias, says Williams, Facebook's global director of diversity. They also must find ways to give those workers more support. To that end, many organizations are turning to people analytics, which aspires to replace gut decisions with data-driven ones. Unfortunately, firms often say that they don't have enough people from marginalized groups in their data sets to produce reliable insights.

But there are things employers can do to supplement small *n*'s: draw on industry or sector data; learn from what's happening in other companies; and deeply examine the experiences of individuals who work for them, talking with them to gather critical qualitative information. If firms are systematic and comprehensive in these efforts, they'll have a better chance of improving diversity and inclusion.

Supplementing the *n*

Nonprofit research organizations are doing important work that sheds light on how bias shapes hiring and advancement in various industries and sectors. For example, a study by the Ascend Foundation showed that in 2013 white men and white women in five major Silicon Valley firms were 154% more likely to become executives than their Asian counterparts were. And though both race and gender were factors in the glass ceiling for Asians, race had 3.7 times the impact that gender did.

It took two more years of research and analysis—using data on several hundred thousand employees, drawn from the EEOC's aggregation of all Bay Area technology firms and from the individual reports of 13 U.S. tech companies—before Ascend determined how bias affected the prospects of blacks and Hispanics. Among those groups it again found that, overall, race had a greater negative impact than gender on advancement from the professional to the executive level. In the Bay Area white women fared worse than white men but much better than all Asians, Hispanics, and blacks. Minority women faced the biggest obstacle to entering the executive ranks. Black and Hispanic women were severely challenged by both their low numbers at the professional level and their lower chances of rising from professional to executive. Asian women, who had more representation at the professional level than other minorities, had the lowest chances of moving up from professional to executive. An analysis of national data found similar results.

By analyzing industry or sector data on underrepresented groups—and examining patterns in hiring, promotions, and other decisions about talent—we can better manage the problems and risks in our own organizations. Tech companies may look at the Ascend reports and say, "Hey, let's think about what's happening with our competitors' talent. There's a good chance it's happening here, too." Their HR teams might then

add a layer of career tracking for women of color, for example, or create training programs for managing diverse teams.

Another approach is to extrapolate lessons from other companies' analyses. We might look, for instance, at Red Ventures, a Charlotte-based digital media company. Red Ventures is diverse by several measures. (It has a Latino CEO, and about 40% of its employees are people of color.) But that doesn't mean there aren't problems to solve. When I met with its top executives, they told me they had recently done an analysis of performance reviews at the firm and found that internalized stereotypes were having a negative effect on black and Latino employees' self-assessments. On average, members of those two groups rated their performance 30% lower than their managers did (whereas white male employees scored their performance 10% higher than their managers did). The study also uncovered a correlation between racial isolation and negative self-perception. For example, people of color who worked in engineering generally rated themselves lower than those who worked in sales, where there were more blacks and Latinos. These patterns were consistent at all levels, from junior to senior staff.

In response, the HR team at Red Ventures trained employees in how to do self-assessments, and that has started to close the gap for blacks and Latinos (who more recently rated themselves 22% lower than their managers did). Hallie Cornetta, the company's VP of human capital, explained that

the training “focused on the importance of completing quantitative and qualitative self-assessments honestly, in a way that shows how employees personally view their performance across our five key dimensions, rather than how they assume their manager or peers view their performance.” She added: “We then shared tangible examples of what ‘exceptional’ versus ‘solid’ versus ‘needs improvement’ looks like in these dimensions to remove some of the subjectivity and help minority—and all— employees assess with greater direction and confidence.”

Getting Personal

Once we’ve gone broader by supplementing the *n,* we can go deeper by examining individual cases. This is critical. Algorithms and statistics do not capture what it feels like to be the only black or Hispanic team member or the effect that marginalization has on individual employees and the group as a whole. We must talk openly with people, one-on-one, to learn about their experiences with bias, and share our own stories to build trust and make the topic safe for discussion. What we discover through those conversations is every bit as important as what shows up in the aggregated data.

An industry colleague, who served as a lead on diversity at a tech company, broke it down for me like this: “When we do our employee surveys, the Latinos always say they are happy. But

I'm Latino, and I know that we are often hesitant to rock the boat. Saying the truth is too risky, so we'll say what you want to hear—even if you sit us down in a focus group. I also know that those aggregated numbers where there are enough of us for the *n* to be significant don't reflect the heterogeneity in our community. Someone who is light-skinned and grew up in Latin America in an upper-middle-class family probably is very happy and comfortable indeed. Someone who is darker-skinned and grew up working-class in America is probably not feeling that same sense of belonging. I'm going to spend time and effort trying to build solutions for the ones I know are at a disadvantage, whether the data tells me that there's a problem with all Latinos or not."

This is a recurring theme. I spoke with 10 diversity and HR professionals at companies with head counts ranging from 60 to 300,000, all of whom are working on programs or interventions for the people who don't register as "big" in big data. They rely at least somewhat on their own intuition when exploring the impact of marginalization. This may seem counter to the mission of people analytics, which is to remove personal perspective and gut feelings from the talent equation entirely. But to discover the effects of bias in our organizations—and to identify complicating factors within groups, such as class and colorism among Latinos and others—we need to collect and analyze qualitative data, too. Intuition can help us find it. The diversity and HR folks described using their "spidey

sense" or knowing there is "something in the water"—essentially, understanding that bias is probably a factor, even though people analytics doesn't always prove causes and predict outcomes. Through conversations with employees—and sometimes through focus groups, if the resources are there and participants feel it's safe to be honest—they reality-check what their instincts tell them, often drawing on their own experiences with bias. One colleague said, "The combination of qualitative and quantitative data is ideal, but at the end of the day there is nothing that data will tell us that we don't already know as black people. I know what my experience was as an African-American man who worked for 16 years in roles that weren't related to improving diversity. It's as much heart as head in this work."

A Call to Action

The proposition at the heart of people analytics is sound—if you want to hire and manage fairly, gut-based decisions are not enough. However, we have to create a new approach, one that also works for small data sets—for the marginalized and the underrepresented.

Here are my recommendations:

First, analysts must challenge the traditional minimum confident *n,* pushing themselves to look beyond the limited hard data. They don't have to prove that the difference in

performance ratings between blacks and whites is “statistically significant” to help managers understand the impact of bias in performance reviews. We already know from the breadth and depth of social science research about bias that it is pervasive in the workplace and influences ratings, so we can combine those insights with what we hear and see on the ground and simply start operating as if bias exists in our companies. We may have to place a higher value on the experiences shared by five or 10 employees—or look more carefully at the descriptive data, such as head counts for underrepresented groups and average job satisfaction scores cut by race and gender—to examine the impact of bias at a more granular level.

In addition, analysts should frequently provide confidence intervals—that is, guidance on how much managers can trust the data if the *n*’s are too small to prove statistical significance. When managers get that information, they’re more likely to make changes in their hiring and management practices, even if they believe—as most do—that they are already treating people fairly. Suppose, for example, that as Red Ventures began collecting data on self-assessments, analysts had a 75% confidence level that blacks and Latinos were underrating themselves. The analysts could then have advised managers to go to their minority direct reports, examine the results from that performance period, and determine together whether the self-reviews truly reflected their contributions. It’s a simple but collaborative way to address implicit bias or stereotyping that

you’re reasonably sure is there while giving agency to each employee.

Second, companies also need to be more consistent and comprehensive in their qualitative analysis. Many already conduct interviews and focus groups to gain insights on the challenges of the underrepresented; some even do textual analysis of written performance reviews, exit interview notes, and hiring memos, looking for language that signals bias or negative stereotyping. But we have to go further. We need to find a viable way to create and process more-objective performance evaluations, given the internalized biases of both employees and managers, and to determine how those biases affect ratings.

This journey begins with educating all employees on the real-life impact of bias and negative stereotypes. At Facebook we offer a variety of training programs with an emphasis on spotting and counteracting bias, and we keep reinforcing key messages post-training, since we know these muscles take time to build. We issue reminders at critical points to shape decision making and behavior. For example, in our performance evaluation tool, we incorporate prompts for people to check word choice when writing reviews and self-assessments. We remind them, for instance, that terms like “cultural fit” can allow bias to creep in and that they should avoid describing women as “bossy” if they wouldn’t describe men who demonstrated the same behaviors that way. We don’t yet have

data on how this is influencing the language used—it's a new intervention—but we will be examining patterns over time.

Perhaps above all, HR and analytics departments must value both qualitative and quantitative expertise and apply mixed-method approaches everywhere possible. At Facebook we're building cross-functional teams with both types of specialists, because no single research method can fully capture the complex layers of bias that everyone brings to the workplace. We view all research methods as trying to solve the same problem from different angles. Sometimes we approach challenges from a quantitative perspective first, to uncover the "what" before looking to the qualitative experts to dive into the "why" and "how." For instance, if the numbers showed that certain teams were losing or attracting minority employees at higher rates than others (the "what"), we might conduct interviews, run focus groups, or analyze text from company surveys to understand the "why," and pull out themes or lessons for other parts of the company. In other scenarios we might reverse the order of those steps. For example, if we repeatedly heard from members of one social group that they weren't seeing their peers getting recognized at the same rate as people in other groups, we could then investigate whether numerical trends confirmed those observations, or conduct statistical analyses to figure out which organizational circumstances were associated with employees' being more or less likely to get recognized.

Cross-functional teams also help us reap the benefits of cognitive diversity. Working together stretches everyone, challenging team members' own assumptions and biases. Getting to absolute "whys" and "hows" on any issue, from recruitment to engagement to performance, is always going to be tough. But we believe that with this approach, we stand the best chance of making improvements across the company. As we analyze the results of Facebook's Pulse survey, given twice a year to employees, and review Performance Summary Cycle inputs, we'll continue to look for signs of problems as well as progress.

Evidence of discrimination or unfair outcomes may not be as certain or obvious in the workplace as it was for me the time I was evicted from my apartment. But we can increase our certainty, and it's essential that we do so. The underrepresented people at our companies are not crazy to perceive biases working against them, and they can get institutional support.

Originally published in November–December 2017. Reprint R1706L

The New CEO Activists

by Aaron K. Chatterji and Michael W. Toffel

WHEN WE FIRST STARTED STUDYING CEO activism, three years ago, we never imagined how significant this phenomenon would become. At the time a small but growing band of executives were taking public stands on political and social issues unrelated to their companies' bottom lines. Since then, controversies over laws affecting transgender people in North Carolina, police shootings in Missouri, and executive orders on immigration have drawn increasing numbers of CEOs into contentious public debates. More recently, the White House's withdrawal from the Paris climate accord, response to the clash between white supremacists and counterprotesters in Charlottesville, Virginia, and decision to rescind Deferred Action for Childhood Arrivals have galvanized many U.S. corporate leaders to speak out and take action.

Of course, corporations have long played an active role in the U.S. political process. They lobby, make contributions to candidates, and fund political action committees and campaigns on various issues in an effort to shape public policies to their

benefit. But CEO activism is something new. Until recently, it was rare for corporate leaders to plunge aggressively into thorny social and political discussions about race, sexual orientation, gender, immigration, and the environment. The so-called Michael Jordan dictum that Republicans buy sneakers too reminds executives that choosing sides on divisive issues can hurt sales, so why do it? Better to weigh in on what traditionally have been seen as business issues, such as taxes and trade, with technocratic arguments rather than moral appeals.

But the world has changed. Political partisanship and discourse grow ever more extreme, and the gridlock in Washington, D.C., shows no sign of easing. Political and social upheaval has provoked frustration and outrage, inspiring business leaders like Tim Cook of Apple, Howard Schultz of Starbucks, and Marc Benioff of Salesforce—among many others—to passionately advocate for a range of causes. “Our jobs as CEOs now include driving what we think is right,” Bank of America’s CEO, Brian Moynihan, told the *Wall Street Journal.* “It’s not exactly political activism, but it is action on issues beyond business.”

The world is taking notice. CEO activism has gotten lots of media attention lately, and public relations firms are now building entire practices around it. While this phenomenon has largely been confined to the United States, there’s little reason to doubt that it could develop into a global force. We believe that the more CEOs speak up on social and political issues, the more they will be expected to do so. And increasingly, CEO activism

has strategic implications: In the Twitter age, silence is more conspicuous—and more consequential.

How CEOs respond: three types of tactics

Traditional	**Nonconfrontational** Lobby behind the scenes Contribute to campaigns Communicate internally with employees Do nothing
Activism	**Raising awareness** Issue a statement or tweet Write an op-ed Seek to spur public action via trade associations
	Exerting economic influence Relocate business activities Pause business expansion Fund political and activist groups

Idea in Brief

The Situation

More and more CEOs are taking a stand on divisive social issues—a dramatic departure from tradition.

The Reason

They’re frustrated with the growing political turmoil and paralysis in the government. Stakeholders, furthermore, are starting to expect corporate leaders to speak out.

The Upshot

CEO activism can have unintended consequences. In this article, the authors look at recent examples of such advocacy and piece together a playbook for executives.

All this activity raises big questions that we will attempt to address: Does CEO activism actually change hearts and minds? What are the risks and potential rewards? And what is the playbook for corporate leaders considering speaking out?

Why CEOs Speak Up

CEOs are weighing in on controversial topics for several reasons. Some point to their corporate values to explain their advocacy, as BOA’s Moynihan and Dan Schulman of PayPal did when taking a stand against a North Carolina law requiring people to use the bathrooms corresponding with the gender on their birth certificates, which became a referendum on transgender rights.

Other CEOs argue that companies should have a higher purpose beyond maximizing shareholder value—a concept that has been gaining traction in the business world. As Benioff told *Time,* “Today CEOs need to stand up not just for their shareholders, but their employees, their customers, their partners, the community, the environment, schools, everybody.”

Activism in action

Corporate leader	Issue	Action taken
Marc Benioff CEO, Salesforce	***Antidiscrimination***	In 2015, Benioff tweeted his opposition to Indiana's Religious Freedom Restoration Act and suspended corporate travel to the state; he later spoke out against North Carolina's bathroom bill and developed a reputation for rallying other business leaders to speak out.
Dan Cathy CEO, Chick-fil-A	***Same-sex marriage***	In 2012, Cathy publicly opposed same-sex marriage on a radio show; his corporation's foundation also donated to anti-LGBTQ organizations.
David and Barbara Green Cofounders, Hobby Lobby	***Health care/ religious freedom***	The Greens filed a highly publicized lawsuit in 2012 to oppose Affordable Care Act–mandated birth control coverage.
Peter Lewis Late chairman, Progressive Insurance	***Marijuana decriminalization***	In 2011, Lewis wrote an opinion piece for Forbes supporting decriminalization; he also donated $3 million to marijuana legalization campaigns.
John Mackey CEO, Whole Foods Market	***Health care***	In 2009, Mackey wrote an editorial criticizing the Affordable Care Act.
Paul Polman CEO, Unilever	***Climate change***	Polman has delivered many public speeches supporting government policies to address climate change.
Jim Rogers Former CEO, Duke Energy	***Climate change***	In 1990, Rogers (as CEO of Public Service Indiana, which eventually became part of Duke Energy) testified before Congress in support of Clean Air Act amendments; he later lobbied Congress to support climate change legislation.
Hamdi Ulukaya CEO, Chobani	***Refugee crisis***	In 2014, Ulukaya pledged to donate $2 million to refugees. He also hired refugees to work at Chobani's manufacturing plants and wrote an op-ed for CNN in support of refugees.

Source: Michael W. Toffel, Aaron K. Chatterji, and Julia Kelley, "CEO Activism (a)," Harvard Business School Case 617-001, March 2017.

And for many leaders, speaking out is a matter of personal conviction. David Green, the founder and CEO of Hobby Lobby, a family-owned chain of crafts stores, cited his religious beliefs when opposing the Obamacare requirement that health insurance for employees include coverage for the morning-after pill among all other forms of birth control.

Some leaders have commented that a greater sense of corporate purpose has become important to Millennials, whether they be employees or customers. Indeed, research from Weber Shandwick and KRC Research finds that large percentages of Millennials believe that CEOs have a responsibility to speak out on political and social issues and say that CEO activism is a factor in their purchasing decisions.

Sometimes leaders point to multiple motivations. "I just think it's insincere to not stand up for those things that you believe in," Jeff Immelt, the former CEO of GE, has said. "We're also stewards of our companies; we're representatives of the people that work with us. And I think we're cowards if we don't take a position occasionally on those things that are really consistent with what our mission is and where our people stand."

The Tactics of CEO Activists

Though they're motivated by diverse interests—external, internal, and deeply personal—activist CEOs generally employ two types of tactics: raising awareness and leveraging economic power.

Raising awareness

For the most part, this involves making public statements—often in the news media, more frequently on Twitter—to garner support for social movements and help usher in change. In such statements business leaders are communicating to stakeholders where they stand on a whole slate of issues that would not have been on the CEO's agenda a generation ago. For example, Goldman Sachs's CEO, Lloyd Blankfein, and Biogen's former CEO George Scangos have spoken out publicly on government policies that affect the rights of LGBTQ individuals. On the socially conservative side of the spectrum, Chick-fil-A's CEO, Dan Cathy, has denounced gay marriage.

In some cases, several CEOs have worked together to raise awareness. For example, days before the United Nations climate-change-agreement negotiations took place in Paris in late 2015, the CEOs of 14 major food companies—Mars, General Mills, Coca-Cola, Unilever, Danone Dairy North America, Hershey, Ben & Jerry's, Kellogg, PepsiCo, Nestlé USA, New Belgium Brewing, Hain Celestial, Stonyfield Farm, and Clif Bar—cosigned an open letter calling on government leaders to create a strong accord that would "meaningfully address the reality of climate change." Similarly, nearly 100 CEOs cosigned an amicus brief to encourage

federal judges to overturn Trump's executive order banning citizens from seven Muslim-majority countries from entering the United States.

Collective action can have greater impact than acting alone. Take what happened with Trump's economic councils. Though Merck's CEO, Kenneth Frazier, received a lot of press when he resigned from the president's American Manufacturing Council in response to Trump's remarks blaming white supremacists and counterprotesters equally for the violence in Charlottesville, it was only after CEOs jumped ship en masse from that group and from Trump's Strategic and Policy Forum that the president disbanded both councils—a move that was widely viewed as a defeat for Trump.

Leveraging economic power

Some of the more powerful cases of CEO activism have involved putting economic pressure on states to reject or overturn legislation. For example, in response to Indiana's Religious Freedom Restoration Act (RFRA), which some viewed as anti-LGBTQ, Bill Oesterle, then the CEO of Angie's List, canceled its planned expansion in Indianapolis, and Benioff threatened to halt all Salesforce employee travel to the state. Other leaders joined the protest, including the president of the National College Athletic Association, Mark Emmert, who suggested that the bill's passage could affect the location of future tournaments and that the association might consider moving its headquarters out of Indianapolis. Under pressure, then-governor Mike Pence

approved a revised version of the law, which forbade businesses from denying service to customers because of their sexual orientation.

In response to North Carolina's bathroom law, Schulman canceled PayPal's plans for a new global operations center in Charlotte, which would have created more than 400 skilled jobs. As many other CEOs followed suit, the potential damage mounted: The Associated Press has estimated that the bathroom law controversy will cost the state more than $3.76 billion in lost business over a dozen years.

Companies and their leaders also wield economic power by donating to third-party groups that promote their favored causes. To help fight Trump's immigration ban, for example, the car-sharing company Lyft pledged $1 million to the American Civil Liberties Union, which is challenging the ban in court. In response to the Charlottesville protest and Trump's reaction to it, James Murdoch, the chief executive of 21st Century Fox, donated $1 million to the Anti-Defamation League, a group that fights bigotry.

How effective are these approaches? The trend of corporate leaders taking a public stand on issues not necessarily related to their businesses is relatively new, so there's little empirical evidence of its impact. But we do have limited anecdotal evidence that it can shape public policy—as it did in the case of Indiana's RFRA. When legislators passed a similar religious freedom bill in Georgia, threats to stop filming in the state from

leaders of many studios and networks—including Disney, CBS, MGM, and Netflix—and similar kinds of warnings from Benioff and other CEOs were seen as instrumental in moving the governor to veto it. And leaders of the National Basketball Association, NCAA, and Atlantic Coast Conference have been credited with forcing North Carolina to revise its bathroom law.

To move beyond anecdotal evidence, we set out to investigate in a scientific, rigorous way whether CEOs can help win public support for policies, thus affecting legislators' votes and whether governors sign or veto bills. Our findings demonstrate that CEOs can indeed play an important role in shaping the public's views on political and social issues. (See the sidebar "Our Research: Does CEO Activism Influence Public Opinion?") Moreover, as we'll discuss, we find that when CEOs communicate a stance on such issues, it can spur like-minded consumers to purchase more of their products.

Our Research: Does CEO Activism Influence Public Opinion?

Some of the experiments we conducted investigated whether and how CEO activism might affect public opinion. In one, we developed a survey asking people if they supported or opposed Indiana's Religious Freedom Restoration Act (RFRA), at a time when the controversy over it was still very much in the news. In some cases, we first told them that many were concerned that the law might allow discrimination against gays and lesbians. In other cases we attributed those concerns to

Apple's CEO, Tim Cook; to Bill Oesterle, who was then CEO of Indiana-based Angie's List; or to the mayor of Indianapolis.

The market research company Civic Science deployed our survey on the hundreds of third-party websites (newspapers, entertainment sites, and so on) with which it partners, gathering 3,418 responses from across the United States. Among those in the baseline condition, who were not told of any discrimination concern, 50% of respondents favored the law—evidence of how split the country is on such legislation. Support for the law dipped to about 40% among respondents who answered the question after being presented with discrimination concerns, regardless of who expressed them—a CEO or a politician—or even if they weren't attributed to anyone in particular.

These results imply that public opinion, at least in this study, was shaped more by the message than by the messenger. There are two ways to interpret this: You can infer that CEOs have no special ability to influence public opinion. After all, their statements had no more effect than politicians' or unattributed statements. On the other hand, the results show that CEOs can be as persuasive as political leaders. CEOs can attract media attention, especially when they speak out on contentious social and environmental issues that are not obviously connected to their bottom lines, which heightens their authenticity. Given that CEOs can sway public opinion, we assume that they can shape public policy, too.

Our study went a bit further to see whether CEO activism would affect people differently depending on their preexisting policy preferences. We found that Cook's discrimination remarks further eroded (already-low) RFRA support among same-sex marriage advocates but had no impact on the much more pro-RFRA views of same-sex marriage opponents. It's important to be aware of whose opinions CEO activism is likely to shift—and whose are likely to be unmoved. In fact, recent research has found that CEOs' political endorsements can significantly affect the campaign contributions of their employees, which suggests that CEO activism might be especially influential with a CEO's own employees.

The Risks and Potential Rewards

In today's politically charged atmosphere, mere affiliations with political leaders or causes can be risky. A few weeks into Trump's term, Under Armour's CEO, Kevin Plank, faced criticism after referring to the president as "a real asset for the country" in an interview. One of his star pitchmen, the Golden State Warriors player Stephen Curry, expressed his displeasure publicly. The hashtag #BoycottUnderArmour began appearing on Twitter, and other Under Armour endorsers, including ballerina Misty Copeland, echoed Curry. The company had to take out a full-page newspaper ad clarifying Plank's comments and stating his opposition to Trump's immigration ban. But that response did not stop Under Armour's stock from being downgraded as one analyst wondered whether the gaffe would "make it nearly impossible to effectively build a cool urban lifestyle brand in the foreseeable future."

CEO activism has sometimes led to charges of hypocrisy. For example, a few conservative websites have criticized Benioff and Cook for denouncing religious freedom laws while Salesforce and Apple continue to do business in countries that persecute LGBTQ individuals. And some activism efforts have come off as clumsy: Consider the widespread ridicule that greeted Howard Schultz's Race Together campaign, in which Starbucks baristas were instructed to write that phrase on all drink cups in an effort to combat racism.

On the other hand, activism can burnish a corporate leader's reputation. In the aftermath of the violence in Charlottesville, the CEOs who resigned from Trump's economic councils (a group that included Plank) were widely praised. The applause for Merck's Frazier, the first to step down, was particularly effusive. "Mr. Frazier, thank you for your courageous stand," tweeted U.S. representative Keith Ellison. The Anne Frank Center for Mutual Respect was even more emphatic, tweeting "A HERO: Ken Frazier."

This controversy also highlighted the risk of silence, which may be viewed as a sign of tacit approval. The *New York Times* and CNBC published lists of which CEOs remained on the president's various economic councils, with CNBC noting that "with each new resignation, those left on the council faced increased scrutiny." Oracle's CEO had similarly been put on the spot when a group of workers from that company launched a petition urging their employer to join numerous other companies in opposing Trump's immigration ban. Their effort attracted national attention, with *USA Today* observing, "More than 130 tech companies—from Apple to Zynga—have signed the amicus brief. Oracle and IBM have not."

Still, CEOs should keep in mind that reactions to activism can cut both ways. While Benioff's advocacy has been widely praised, he admitted to CBS News that Colin Powell, the former secretary of state and a retired four-star general—and now a Salesforce director—warned him: "The farther you go up the

tree, the more your backside is going to be exposed, and you'd better be careful." After Chick-fil-A's Cathy spoke out against gay marriage, the chain faced consumer picket lines and a boycott—but also a countervailing "Chick-fil-A Appreciation Day," which attracted large crowds of customers. Indeed, in a Weber Shandwick survey 40% of respondents said they would be more likely to purchase from a company if they agreed with the CEO's position, but 45% said they'd be less likely to if they disagreed with the CEO's view.

We conducted our own experiment to assess the influence of CEO activism on U.S. consumers' behavior. In it, we asked a nationally representative group of respondents about their intent to buy Apple products in the near future. To some, we first provided a statement describing CEO Tim Cook's opinion that Indiana's religious freedom bill was discriminatory against LGBTQ individuals; to others, we provided a generic statement about Cook's management philosophy. To the rest, we provided no statement at all; we simply asked about purchasing intent. We randomly deployed these three conditions and received 2,176 responses. The people in the group exposed to Cook's activism, we found, expressed significantly higher intent to buy Apple products in the near future than those in the other two groups. Learning about Cook's activism increased intent to purchase among supporters of same-sex marriage but did not erode intent among its opponents. These results indicate that CEO activism can generate goodwill for the company but need not alienate

those who disagree with the CEO. But this most likely does not apply to all companies. Apple products are especially sticky, so while Cook's remarks might not provoke a backlash against iPhones, other business leaders should consider whether the political makeup of their consumers and the nature of their products might lead to a different result. It's critical for every CEO to proceed thoughtfully.

The CEO Activist's Playbook

Drawing on our empirical research and interviews with CEO activists and their stakeholders, we have developed a guide for leaders who are deciding whether to speak out and how.

What to weigh in on

Smart CEO activists typically choose their issues; the issues do not choose them. To avoid being blindsided by a news story or awkwardly weighing in on a topic they know little about, CEOs should sit down with their executive teams, including their chief communications officers, and decide what issues matter to them and why. This discussion should include reflection on why championing the selected causes would have greater social impact than championing other causes. (On occasion, however, there's no time for this kind of deliberation, such as when corporate leaders felt they quickly needed to make it clear they had no tolerance for racism after Charlottesville.)

Executives must balance the likelihood of having an effect and other potential benefits—such as pleasing employees and consumers—against the possibility of a backlash. As part of this assessment, CEOs should explicitly consider how their statements and actions will be received in a politically polarized atmosphere. A 2016 Global Strategy Group report shows that when companies are associated with political issues, customers view this connection through the lens of their party affiliation. (See the exhibit "A polarized response.") According to the study, twice as many Democrats viewed Schultz's Race Together campaign positively as viewed it negatively, but three times as many Republicans viewed it unfavorably as viewed it favorably. Cook's advocacy for gay marriage produced similar responses. Championship of less divisive issues, such as parental leave and STEM education, however, is more likely to improve the brand image of the CEO's company among both Democrats and Republicans, the study showed.

A polarized response

Democrats and Republicans can have very different reactions to corporate activism.

The chart below shows how each company's stance on a social issue affected its overall favorability ratings with Democrats and Republicans. The percentages indicate the net change in support from members of each party in response to the activist stance.

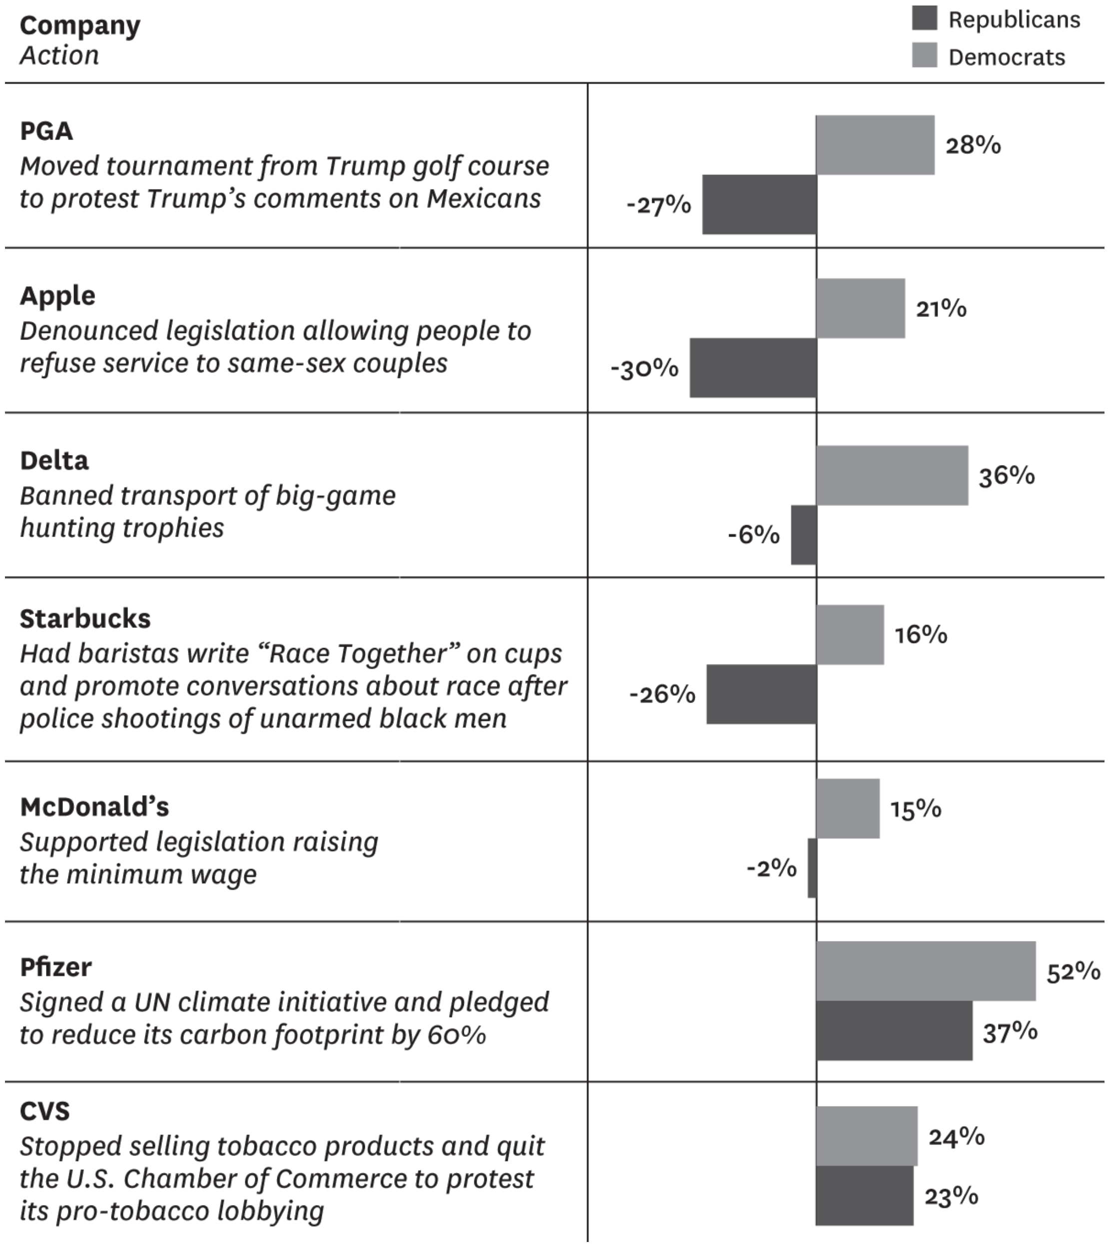

Source: "Business & Politics: Do They Mix?" Third Annual Study, January 2016, a survey of 803 U.S. adults by Global Strategy Group.

CEOs should also consider the extent to which the public believes a CEO voice is appropriate on a given topic. The Global Strategy Group study found that Democrats and Republicans both thought it was fitting for companies to take public stances

on economic issues like minimum wage and parental leave. However, there was much less consensus about the appropriateness of weighing in on social issues such as abortion, gun control, LGBTQ equality, and immigration. (See the exhibit "Is it appropriate to take a stand? What consumers think.")

Is it appropriate to take a stand? What consumers think

A Global Strategy Group survey showed that Americans tend to approve of corporate activism on economic issues more than activism on social issues.

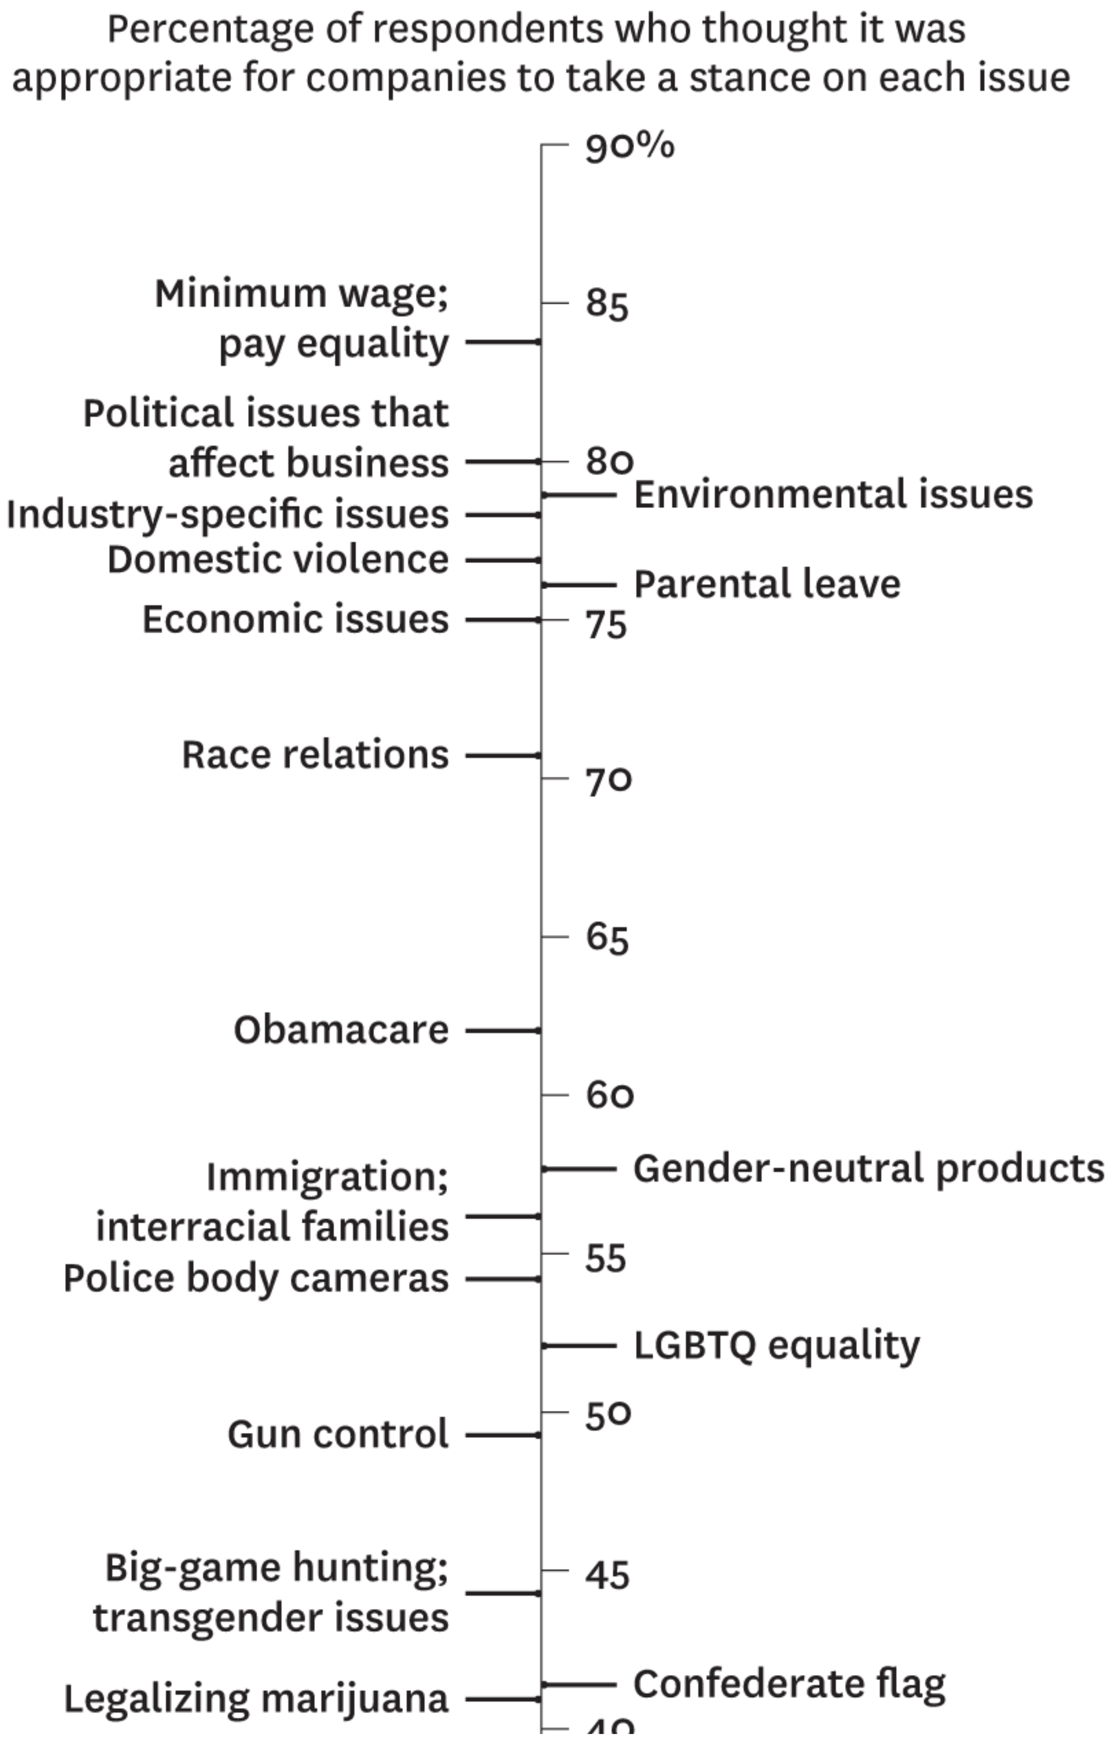
Percentage of respondents who thought it was appropriate for companies to take a stance on each issue
90%
85
80
75
70
65
60
55
50
45
Minimum wage; pay equality
Political issues that affect business
Environmental issues
Industry-specific issues
Domestic violence
Parental leave
Economic issues
Race relations
Obamacare
Gender-neutral products
Immigration; interracial families
Police body cameras
LGBTQ equality
Gun control
Big-game hunting; transgender issues
Confederate flag
Legalizing marijuana

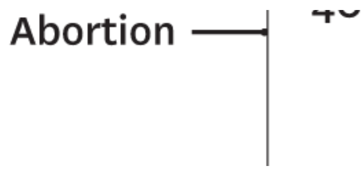

Source: "Business & Politics: Do They Mix?" Third Annual Study, January 2016, a survey of 803 U.S. Adults by Global Strategy Group.

Immigration has proven a particularly complex issue, as the experiences of Chobani's CEO, Hamdi Ulukaya, and Carbonite's CEO, Mohamad Ali, illustrate. Immigrants to the United States themselves, both publicly opposed the Trump administration's restrictions. Both have been praised for their stances, but Ulukaya was also threatened and his company faced a boycott, while Ali's remarks prompted no discernible backlash. This difference could be attributed to Ulukaya's focus on the moral need t o provide job opportunities for refugees, whereas Ali placed more emphasis on immigrants as job creators whose work also benefits native-born citizens. It's important to note, however, that while speaking out on controversial topics might provoke an adverse reaction, it is also likely to attract media coverage, which increases the opportunity for a CEO's views to be heard in the first place.

To influence public policy, the message has to be authentic to both the individual leader and the business. There should be a compelling narrative for why *this* issue matters to *this* CEO of *this* business at *this* time. The issue selection is also a crucial time to "get smart" about the underlying details. CEOs can quickly get in over their heads if they start speaking publicly about complex

issues and are pressed by knowledgeable journalists and commentators. Because the credibility of business leaders rests on the perception that they make decisions after careful analysis, CEO activists can be effective only if they really understand the issue under debate.

When to weigh in

Once the issue is selected, the CEO activist has to understand if there are key moments when speaking out might actually make a difference. Is it while a piece of legislation is being considered, or is it afterward?

We have observed that a CEO activist's chances of blocking a particular policy are typically better than his or her chances of reversing legislation that has been enacted. As we have seen with the Republican Party's efforts to repeal the Affordable Care Act in recent months, the U.S. legislative system was designed to be slow moving and deliberative. This institutional feature makes it difficult not only to pass sweeping new legislation but to repeal existing laws as well.

Also, consider the news cycle. As we noted earlier, being the first CEO to quit one of the president's economic councils earned Frazier (and Merck) significant positive media coverage. When other CEOs quit in rapid succession over the next 48 hours, their stories were lumped together. Frazier's actions will likely be remembered more than those of the CEOs who followed him. Of course, there was a downside to all the attention: President Trump struck back directly at Frazier, tweeting an insult and

citing Merck's responsibility for high drug prices. To date, there's no evidence that this has hurt Merck's business.

Implications for Democracy

CEO activism may be giving businesses and their leaders even more influence in a political system in which their money can already buy access to power. Some people, including North Carolina's lieutenant governor, who supported the bathroom bill while facing an onslaught of CEO activism, have gone further, characterizing it as corporate bullying. One Georgia state senator, who sponsored that state's religious freedom bill, lamented, "Marc Benioff is the ringleader for big-business CEOs who use economic threats to exercise more power over public policy than the voters who use the democratic process." From this perspective, CEO activism can be viewed as endangering democracy's ideal that each citizen should have an equal say in influencing policy outcomes.

There is of course another angle on this that considers CEO activism within the current environment of political influence. As we've noted, CEO activism is an unusually transparent way for corporate leaders to try to affect policy—in contrast to behind-the-scenes efforts to work with legislators, trade associations, and think tanks. Because CEO activism is highly visible, employees, customers, and the media can decide how to respond to it. There is also a political divide here. (To be sure, certain controversies transcend politics.) Some progressives have been appreciative of recent CEO activism while decrying the activities of business leaders like the Koch brothers. As a result, many conservatives see a double standard at play. Most of the CEO activists have been espousing liberal views, but it remains to be seen how widespread activism from conservative business leaders would be received.

How to weigh in

CEO activism differs from traditional corporate engagement in politics precisely because it is visible and high profile. The CEO needs to decide whether he or she wants all that attention or if the cause would be better advanced by a coalition of CEOs. More than 160 CEOs and business leaders chose to sign a letter by the Human Rights Campaign opposing the North Carolina bathroom law. In taking this approach, they mitigated the risk of consumer backlash and amplified the newsworthiness and thus the impact of their activism. Collective action can also make it more difficult for critics to target individual corporate leaders and thus can be perceived as less risky. But it is slower by design and is likely to be less effective in associating a particular leader and corporate brand with a particular cause.

CEOs also may choose not to weigh in at all. Some leaders may feel that they do not understand the issue well enough, hold an unpopular view, or simply want to focus on other areas. All of those are credible reasons to hold back. But executives should expect that employees, the media, and other interested parties may ask why the CEO has not spoken out, and should be ready to explain the rationale.

The inside game

It's a good idea to make sure that internal stakeholders are aligned with CEO activism—or at least aware of it ahead of time. When Frazier was considering resigning from Trump's economic council, he reached out to his board members, who subsequently defended his decision and praised his courage and integrity. Our

interviews suggest that not all CEOs consult with their directors or employees before taking public stands, which may imperil their efforts.

Though CEOs first have to decide whether they're speaking for themselves or their organizations, they should recognize that any statements they make will nonetheless be associated with their companies. We have seen almost no CEOs successfully separate themselves from their firms in this way. Given that, we advise setting up a rapid response team composed of representatives from the board, investors, senior management (including the chief communications officer), and employees to act as a kitchen cabinet on CEO activism. Seeking broad consensus across the organization could prevent CEO activism from being timely, which is often critical to attract attention to a message, but if the CEO can at least inform his or her cabinet about what to expect and why, it should greatly reduce the risk that key stakeholders will be unprepared for any backlash.

Predicting the reaction and gauging the results

CEO activists should prepare thoughtful responses to those who disagree with them. After Target modified its bathroom policy to accommodate transgender customers, hundreds of thousands of people signed a petition in protest. The literature tells us that when easy alternatives to a product or service are available, boycotts are more effective. Target is particularly vulnerable in this regard. Thus it's not surprising that the retail chain, which has many stores in politically conservative areas of the United

States, has taken action to assuage the criticism by spending $20 million creating single-occupancy bathrooms in its stores. On the other hand, Nordstrom's customer base of affluent urban women did not threaten to abandon the upscale department store chain when President Trump attacked it for distancing itself from Ivanka Trump's apparel line.

Companies generally lack good data on the political beliefs of their customers, but this information would be useful in assessing potential reactions to CEO activism. CEOs and their companies are likely to know more about the political beliefs of their employees and can better predict their responses, however. Will employees rally to the cause or go public with their disapproval—as more than a thousand IBM employees did after CEO Virginia Rometty met with President Trump?

CEO activism also risks a backlash from politicians. Trump has tweeted his disagreement with numerous companies and their management decisions, marshaling millions of Twitter followers and creating public relations headaches. CEOs and their teams should be gaming out the likely response from supporters and critics in their own organizations, the media, and the political sphere.

It's imperative to hold postmortems, too, and answer the question: Did I make a difference? Metrics to assess the impact of activism should be established ahead of time, whether they be retweets, media mentions, public opinion polls, or actual policy shifts. Big swings in public opinion are rare, so it makes sense to

set realistic goals, track intermediate outcomes, and measure progress over time.

CEO activism could become a first-order strategic issue. As more and more business leaders choose to speak out on contentious political and social matters, CEOs will increasingly be called on to help shape the debate about such issues. Many will decide to stay out of the fray, but they should still expect to be peppered with questions from employees, the media, and other stakeholders about the hot-button topics of day.

We believe CEOs need a playbook in this new world. To effectively engage in CEO activism, they should select issues carefully, reflect on the best times and approaches to get involved, consider the potential for backlash, and measure results. By following these guidelines, CEO activists can be more effective on the issues they care about most.

Further Reading

- **"Why Apple's Tim Cook and Other CEOs Are Speaking Out on Police Shootings,"** Aaron K. Chatterji, *Fortune*, July 16, 2016
- **"Do CEO Activists Make a Difference? Evidence from a Quasi-Field Experiment,"** Aaron K. Chatterji and Michael W. Toffel, working paper, July 2017

- **“Starbucks’ ‘Race Together’ Campaign and the Upside of CEO Activism,”** Aaron K. Chatterji and Michael W. Toffel, *Harvard Business Review*, March 24, 2015
- **“The Power of C.E.O. Activism,”** Aaron K. Chatterji and Michael W. Toffel, *New York Times*, April 1, 2016
- **“Is It Safe for CEOs to Voice Strong Political Opinions?”** Leslie Gaines-Ross, *Harvard Business Review*, June 23, 2016
- **“Business & Politics: Do They Mix?”** Global Strategy Group, annual studies, 2013–2016
- **“The Dawn of CEO Activism,”** Weber Shandwick, with KRC Research, 2016

Originally published in January–February 2018. Reprint R1801E

Artificial Intelligence for the Real World

by Thomas H. Davenport and Rajeev Ronanki

IN 2013, THE MD ANDERSON CANCER CENTER launched a "moon shot" project: diagnose and recommend treatment plans for certain forms of cancer using IBM's Watson cognitive system. But in 2017, the project was put on hold after costs topped $62 million—and the system had yet to be used on patients. At the same time, the cancer center's IT group was experimenting with using cognitive technologies to do much less ambitious jobs, such as making hotel and restaurant recommendations for patients' families, determining which patients needed help paying bills, and addressing staff IT problems. The results of these projects have been much more promising: The new systems have contributed to increased patient satisfaction, improved financial performance, and a decline in time spent on tedious data entry by the hospital's care managers. Despite the setback on the moon shot, MD Anderson remains committed to using cognitive technology—that is, next-generation artificial intelligence—to enhance cancer treatment, and is currently developing a variety

of new projects at its center of competency for cognitive computing.

The contrast between the two approaches is relevant to anyone planning AI initiatives. Our survey of 250 executives who are familiar with their companies' use of cognitive technology shows that three-quarters of them believe that AI will substantially transform their companies within three years. However, our study of 152 projects in almost as many companies also reveals that highly ambitious moon shots are less likely to be successful than "low-hanging fruit" projects that enhance business processes. This shouldn't be surprising—such has been the case with the great majority of new technologies that companies have adopted in the past. But the hype surrounding artificial intelligence has been especially powerful, and some organizations have been seduced by it.

In this article, we'll look at the various categories of AI being employed and provide a framework for how companies should begin to build up their cognitive capabilities in the next several years to achieve their business objectives.

Three Types of AI

It is useful for companies to look at AI through the lens of business capabilities rather than technologies. Broadly speaking, AI can support three important business needs: automating business processes, gaining insight through data analysis, and

engaging with customers and employees. (See the exhibit “Cognitive projects by type.”)

Cognitive projects by type

We studied 152 cognitive technology projects and found that they fell into three categories.

Robotics & cognitive automation:	Cognitive insight:	Cognitive engagement:
71	57	24

Idea in Brief

The Problem

Cognitive technologies are increasingly being used to solve business problems, but many of the most ambitious AI projects encounter setbacks or fail.

The Approach

Companies should take an incremental rather than a transformative approach and focus on augmenting rather than replacing human capabilities.

The Process

To get the most out of AI, firms must understand which technologies perform what types of tasks, create a prioritized portfolio of projects based on business needs, and develop plans to scale up across the company.

Process automation

Of the 152 projects we studied, the most common type was the automation of digital and physical tasks—typically back-office administrative and financial activities—using robotic process automation technologies. RPA is more advanced than earlier business-process automation tools, because the "robots" (that is, code on a server) act like a human inputting and consuming information from multiple IT systems. Tasks include:

- transferring data from e-mail and call center systems into systems of record—for example, updating customer files with address changes or service additions;
- replacing lost credit or ATM cards, reaching into multiple systems to update records and handle customer communications;
- reconciling failures to charge for services across billing systems by extracting information from multiple document types; and
- "reading" legal and contractual documents to extract provisions using natural language processing.

RPA is the least expensive and easiest to implement of the cognitive technologies we'll discuss here, and typically brings a quick and high return on investment. (It's also the least "smart" in the sense that these applications aren't programmed to learn and improve, though developers are slowly adding more

intelligence and learning capability.) It is particularly well suited to working across multiple back-end systems.

The business benefits of AI

We surveyed 250 executives who were familiar with their companies' use of cognitive technologies to learn about their goals for AI initiatives. More than half said their primary goal was to make existing products better. Reducing head count was mentioned by only 22%.

Source: Deloitte 2017.

At NASA, cost pressures led the agency to launch four RPA pilots in accounts payable and receivable, IT spending, and

human resources—all managed by a shared services center. The four projects worked well—in the HR application, for example, 86% of transactions were completed without human intervention—and are being rolled out across the organization. NASA is now implementing more RPA bots, some with higher levels of intelligence. As Jim Walker, project leader for the shared services organization notes, “So far it’s not rocket science.”

The challenges of AI

Executives in our survey identified several factors that can stall or derail AI initiatives, ranging from integration issues to scarcity of talent.

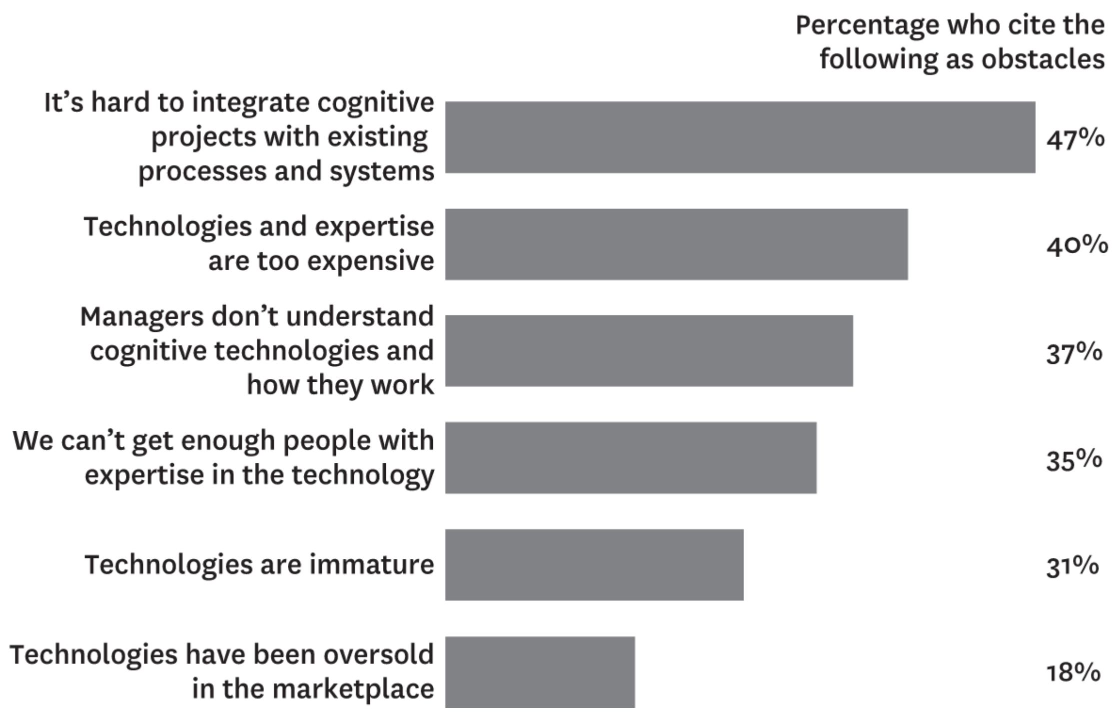

Source: Deloitte 2017.

One might imagine that robotic process automation would quickly put people out of work. But across the 71 RPA projects we reviewed (47% of the total), replacing administrative employees was neither the primary objective nor a common outcome. Only a few projects led to reductions in head count, and in most cases, the tasks in question had already been shifted to outsourced workers. As technology improves, robotic automation projects are likely to lead to some job losses in the future, particularly in the offshore business-process outsourcing industry. If you can outsource a task, you can probably automate it.

Cognitive insight

The second most common type of project in our study (38% of the total) used algorithms to detect patterns in vast volumes of data and interpret their meaning. Think of it as "analytics on steroids." These machine-learning applications are being used to:

- predict what a particular customer is likely to buy;
- identify credit fraud in real time and detect insurance claims fraud;
- analyze warranty data to identify safety or quality problems in automobiles and other manufactured products;
- automate personalized targeting of digital ads; and
- provide insurers with more-accurate and detailed actuarial modeling.

Cognitive insights provided by machine learning differ from those available from traditional analytics in three ways: They are usually much more data-intensive and detailed, the models typically are trained on some part of the data set, and the models get better—that is, their ability to use new data to make predictions or put things into categories improves over time.

Versions of machine learning (deep learning, in particular, which attempts to mimic the activity in the human brain in order to recognize patterns) can perform feats such as recognizing images and speech. Machine learning can also make available new data for better analytics. While the activity of data curation has historically been quite labor-intensive, now machine learning can identify probabilistic matches—data that is likely to be associated with the same person or company but that appears in slightly different formats—across databases. GE has used this technology to integrate supplier data and has saved $80 million in its first year by eliminating redundancies and negotiating contracts that were previously managed at the business unit level. Similarly, a large bank used this technology to extract data on terms from supplier contracts and match it with invoice numbers, identifying tens of millions of dollars in products and services not supplied. Deloitte's audit practice is using cognitive insight to extract terms from contracts, which enables an audit to address a much higher proportion of documents, often 100%, without human auditors' having to painstakingly read through them.

Cognitive insight applications are typically used to improve performance on jobs only machines can do—tasks such as programmatic ad buying that involve such high-speed data crunching and automation that they've long been beyond human ability—so they're not generally a threat to human jobs.

Cognitive engagement

Projects that engage employees and customers using natural language processing chatbots, intelligent agents, and machine learning were the least common type in our study (accounting for 16% of the total). This category includes:

- intelligent agents that offer 24/7 customer service addressing a broad and growing array of issues from password requests to technical support questions—all in the customer's natural language;
- internal sites for answering employee questions on topics including IT, employee benefits, and HR policy;
- product and service recommendation systems for retailers that increase personalization, engagement, and sales—typically including rich language or images; and
- health treatment recommendation systems that help providers create customized care plans that take into account individual patients' health status and previous treatments.

The companies in our study tended to use cognitive engagement technologies more to interact with employees than with customers. That may change as firms become more comfortable turning customer interactions over to machines. Vanguard, for example, is piloting an intelligent agent that helps its customer service staff answer frequently asked questions. The plan is to eventually allow customers to engage with the cognitive agent directly, rather than with the human customer-service agents. SEBank, in Sweden, and the medical technology giant Becton, Dickinson, in the United States, are using the lifelike intelligent-agent avatar Amelia to serve as an internal employee help desk for IT support. SEBank has recently made Amelia available to customers on a limited basis in order to test its performance and customer response.

Companies tend to take a conservative approach to customer-facing cognitive engagement technologies largely because of their immaturity. Facebook, for example, found that its Messenger chatbots couldn't answer 70% of customer requests without human intervention. As a result, Facebook and several other firms are restricting bot-based interfaces to certain topic domains or conversation types.

Our research suggests that cognitive engagement apps are not currently threatening customer service or sales rep jobs. In most of the projects we studied, the goal was not to reduce head count but to handle growing numbers of employee and customer interactions without adding staff. Some organizations were

planning to hand over routine communications to machines, while transitioning customer-support personnel to more-complex activities such as handling customer issues that escalate, conducting extended unstructured dialogues, or reaching out to customers before they call in with problems.

As companies become more familiar with cognitive tools, they are experimenting with projects that combine elements from all three categories to reap the benefits of AI. An Italian insurer, for example, developed a "cognitive help desk" within its IT organization. The system engages with employees using deep-learning technology (part of the cognitive insights category) to search frequently asked questions and answers, previously resolved cases, and documentation to come up with solutions to employees' problems. It uses a smart-routing capability (business process automation) to forward the most complex problems to human representatives, and it uses natural language processing to support user requests in Italian.

Despite their rapidly expanding experience with cognitive tools, however, companies face significant obstacles in development and implementation. On the basis of our research, we've developed a four-step framework for integrating AI technologies that can help companies achieve their objectives, whether the projects are moon shoots or business-process enhancements.

1. Understanding the Technologies

Before embarking on an AI initiative, companies must understand which technologies perform what types of tasks, and the strengths and limitations of each. Rule-based expert systems and robotic process automation, for example, are transparent in how they do their work, but neither is capable of learning and improving. Deep learning, on the other hand, is great at learning from large volumes of labeled data, but it's almost impossible to understand how it creates the models it does. This "black box" issue can be problematic in highly regulated industries such as financial services, in which regulators insist on knowing why decisions are made in a certain way.

We encountered several organizations that wasted time and money pursuing the wrong technology for the job at hand. But if they're armed with a good understanding of the different technologies, companies are better positioned to determine which might best address specific needs, which vendors to work with, and how quickly a system can be implemented. Acquiring this understanding requires ongoing research and education, usually within IT or an innovation group.

In particular, companies will need to leverage the capabilities of key employees, such as data scientists, who have the statistical and big-data skills necessary to learn the nuts and bolts of these technologies. A main success factor is your people's willingness to learn. Some will leap at the opportunity, while others will want to stick with tools they're familiar with. Strive to have a high percentage of the former.

If you don't have data science or analytics capabilities in-house, you'll probably have to build an ecosystem of external service providers in the near term. If you expect to be implementing longer-term AI projects, you will want to recruit expert in-house talent. Either way, having the right capabilities is essential to progress.

Given the scarcity of cognitive technology talent, most organizations should establish a pool of resources—perhaps in a centralized function such as IT or strategy—and make experts available to high-priority projects throughout the organization. As needs and talent proliferate, it may make sense to dedicate groups to particular business functions or units, but even then a central coordinating function can be useful in managing projects and careers.

2. Creating a Portfolio of Projects

The next step in launching an AI program is to systematically evaluate needs and capabilities and then develop a prioritized portfolio of projects. In the companies we studied, this was usually done in workshops or through small consulting engagements. We recommend that companies conduct assessments in three broad areas.

Identifying the opportunities

The first assessment determines which areas of the business could benefit most from cognitive applications. Typically, they

are parts of the company where “knowledge”—insight derived from data analysis or a collection of texts—is at a premium but for some reason is not available.

- **Bottlenecks.** In some cases, the lack of cognitive insights is caused by a bottleneck in the flow of information; knowledge exists in the organization, but it is not optimally distributed. That’s often the case in health care, for example, where knowledge tends to be siloed within practices, departments, or academic medical centers.
- **Scaling challenges.** In other cases, knowledge exists, but the process for using it takes too long or is expensive to scale. Such is often the case with knowledge developed by financial advisers. That’s why many investment and wealth management firms now offer AI-supported “robo-advice” capabilities that provide clients with cost-effective guidance for routine financial issues.

  In the pharmaceutical industry, Pfizer is tackling the scaling problem by using IBM’s Watson to accelerate the laborious process of drug-discovery research in immuno-oncology, an emerging approach to cancer treatment that uses the body’s immune system to help fight cancer. Immuno-oncology drugs can take up to 12 years to bring to market. By combining a sweeping literature review with Pfizer’s own data, such as lab reports, Watson is helping researchers to surface relationships and find hidden

patterns that should speed the identification of new drug targets, combination therapies for study, and patient selection strategies for this new class of drugs.

- **Inadequate firepower.** Finally, a company may collect more data than its existing human or computer firepower can adequately analyze and apply. For example, a company may have massive amounts of data on consumers' digital behavior but lack insight about what it means or how it can be strategically applied. To address this, companies are using machine learning to support tasks such as programmatic buying of personalized digital ads or, in the case of Cisco Systems and IBM, to create tens of thousands of "propensity models" for determining which customers are likely to buy which products.

Determining the use cases

The second area of assessment evaluates the use cases in which cognitive applications would generate substantial value and contribute to business success. Start by asking key questions such as: How critical to your overall strategy is addressing the targeted problem? How difficult would it be to implement the proposed AI solution—both technically and organizationally? Would the benefits from launching the application be worth the effort? Next, prioritize the use cases according to which offer the most short- and long-term value, and which might ultimately be

integrated into a broader platform or suite of cognitive capabilities to create competitive advantage.

Selecting the technology

The third area to assess examines whether the AI tools being considered for each use case are truly up to the task. Chatbots and intelligent agents, for example, may frustrate some companies because most of them can't yet match human problem solving beyond simple scripted cases (though they are improving rapidly). Other technologies, like robotic process automation that can streamline simple processes such as invoicing, may in fact slow down more-complex production systems. And while deep learning visual recognition systems can recognize images in photos and videos, they require lots of labeled data and may be unable to make sense of a complex visual field.

In time, cognitive technologies will transform how companies do business. Today, however, it's wiser to take incremental steps with the currently available technology while planning for transformational change in the not-too-distant future. You may ultimately want to turn customer interactions over to bots, for example, but for now it's probably more feasible—and sensible—to automate your internal IT help desk as a step toward the ultimate goal.

3. Launching Pilots

Because the gap between current and desired AI capabilities is not always obvious, companies should create pilot projects for cognitive applications before rolling them out across the entire enterprise.

Proof-of-concept pilots are particularly suited to initiatives that have high potential business value or allow the organization to test different technologies at the same time. Take special care to avoid "injections" of projects by senior executives who have been influenced by technology vendors. Just because executives and boards of directors may feel pressure to "do something cognitive" doesn't mean you should bypass the rigorous piloting process. Injected projects often fail, which can significantly set back the organization's AI program.

If your firm plans to launch several pilots, consider creating a cognitive center of excellence or similar structure to manage them. This approach helps build the needed technology skills and capabilities within the organization, while also helping to move small pilots into broader applications that will have a greater impact. Pfizer has more than 60 projects across the company that employ some form of cognitive technology; many are pilots, and some are now in production.

At Becton, Dickinson, a "global automation" function within the IT organization oversees a number of cognitive technology pilots that use intelligent digital agents and RPA (some work is done in partnership with the company's Global Shared Services organization). The global automation group uses end-to-end

process maps to guide implementation and identify automation opportunities. The group also uses graphical “heat maps” that indicate the organizational activities most amenable to AI interventions. The company has successfully implemented intelligent agents in IT support processes, but as yet is not ready to support large-scale enterprise processes, like order-to-cash. The health insurer Anthem has developed a similar centralized AI function that it calls the Cognitive Capability Office.

Business-process redesign

As cognitive technology projects are developed, think through how workflows might be redesigned, focusing specifically on the division of labor between humans and the AI. In some cognitive projects, 80% of decisions will be made by machines and 20% will be made by humans; others will have the opposite ratio. Systematic redesign of workflows is necessary to ensure that humans and machines augment each other’s strengths and compensate for weaknesses.

The investment firm Vanguard, for example, has a new “Personal Advisor Services” (PAS) offering, which combines automated investment advice with guidance from human advisers. In the new system, cognitive technology is used to perform many of the traditional tasks of investment advising, including constructing a customized portfolio, rebalancing portfolios over time, tax loss harvesting, and tax-efficient investment selection. Vanguard’s human advisers serve as “investing coaches,” tasked with answering investor questions,

encouraging healthy financial behaviors, and being, in Vanguard's words, "emotional circuit breakers" to keep investors on plan. Advisers are encouraged to learn about behavioral finance to perform these roles effectively. The PAS approach has quickly gathered more than $80 billion in assets under management, costs are lower than those for purely human-based advising, and customer satisfaction is high. (See the exhibit "One company's division of labor.")

One company's division of labor

Vanguard, the investment services firm, uses cognitive technology to provide customers with investment advice at a lower cost. Its Personal Advisor Services system automates many traditional tasks of investment advising, while human advisers take on higher-value activities. Here's how Vanguard redesigned its work processes to get the most from the new system.

Cognitive technology	Adviser
Generates a financial plan	Understands investment goals
Provides goals-based forecasting in real time	Customizes an implementation plan
Rebalances portfolio to target mix	Provides investment analysis and retirement planning
Minimizes taxes	Develops retirement income and Social Security drawdown strategies
Tracks aggregated assets in one place	Serves as a behavioral coach
Engages clients virtually	Monitors spending to encourage accountability
	Offers ongoing wealth and financial-planning support
	Addresses estate-planning considerations

Source: Vanguard Group.

Vanguard understood the importance of work redesign when implementing PAS, but many companies simply “pave the cow path” by automating existing work processes, particularly when using RPA technology. By automating established workflows, companies can quickly implement projects and achieve ROI—but they forgo the opportunity to take full advantage of AI capabilities and substantively improve the process.

Cognitive work redesign efforts often benefit from applying design-thinking principles: understanding customer or end-user needs, involving employees whose work will be restructured, treating designs as experimental “first drafts,” considering multiple alternatives, and explicitly considering cognitive technology capabilities in the design process. Most cognitive projects are also suited to iterative, agile approaches to development.

4. Scaling Up

Many organizations have successfully launched cognitive pilots, but they haven’t had as much success rolling them out organization-wide. To achieve their goals, companies need detailed plans for scaling up, which requires collaboration between technology experts and owners of the business process being automated. Because cognitive technologies typically support individual tasks rather than entire processes, scale-up almost always requires integration with existing systems and processes. Indeed, in our survey, executives reported that such integration was the greatest challenge they faced in AI initiatives.

Companies should begin the scaling-up process by considering whether the required integration is even possible or feasible. If the application depends on special technology that is difficult to source, for example, that will limit scale-up. Make sure your business process owners discuss scaling considerations with the

IT organization before or during the pilot phase: An end run around IT is unlikely to be successful, even for relatively simple technologies like RPA.

The health insurer Anthem, for example, is taking on the development of cognitive technologies as part of a major modernization of its existing systems. Rather than bolting new cognitive apps onto legacy technology, Anthem is using a holistic approach that maximizes the value being generated by the cognitive applications, reduces the overall cost of development and integration, and creates a halo effect on legacy systems. The company is also redesigning processes at the same time to, as CIO Tom Miller puts it, "use cognitive to move us to the next level."

In scaling up, companies may face substantial change-management challenges. At one U.S. apparel retail chain, for example, the pilot project at a small subset of stores used machine learning for online product recommendations, predictions for optimal inventory and rapid replenishment models, and—most difficult of all—merchandising. Buyers, used to ordering product on the basis of their intuition, felt threatened and made comments like "If you're going to trust this, what do you need me for?" After the pilot, the buyers went as a group to the chief merchandising officer and requested that the program be killed. The executive pointed out that the results were positive and warranted expanding the project. He assured the buyers that, freed of certain merchandising tasks, they could take on more high-value work that humans can still do better

than machines, such as understanding younger customers' desires and determining apparel manufacturers' future plans. At the same time, he acknowledged that the merchandisers needed to be educated about a new way of working.

If scale-up is to achieve the desired results, firms must also focus on improving productivity. Many, for example, plan to grow their way into productivity—adding customers and transactions without adding staff. Companies that cite head count reduction as the primary justification for the AI investment should ideally plan to realize that goal over time through attrition or from the elimination of outsourcing.

The Future Cognitive Company

Our survey and interviews suggest that managers experienced with cognitive technology are bullish on its prospects. Although the early successes are relatively modest, we anticipate that these technologies will eventually transform work. We believe that companies that are adopting AI in moderation now—and have aggressive implementation plans for the future—will find themselves as well positioned to reap benefits as those that embraced analytics early on.

Through the application of AI, information-intensive domains such as marketing, health care, financial services, education, and professional services could become simultaneously more valuable and less expensive to society. Business drudgery in every industry and function—overseeing routine transactions,

repeatedly answering the same questions, and extracting data from endless documents—could become the province of machines, freeing up human workers to be more productive and creative. Cognitive technologies are also a catalyst for making other data-intensive technologies succeed, including autonomous vehicles, the Internet of Things, and mobile and multichannel consumer technologies.

The great fear about cognitive technologies is that they will put masses of people out of work. Of course, some job loss is likely as smart machines take over certain tasks traditionally done by humans. However, we believe that most workers have little to fear at this point. Cognitive systems perform tasks, not entire jobs. The human job losses we've seen were primarily due to attrition of workers who were not replaced or through automation of outsourced work. Most cognitive tasks currently being performed augment human activity, perform a narrow task within a much broader job, or do work that wasn't done by humans in the first place, such as big-data analytics.

Most managers with whom we discuss the issue of job loss are committed to an augmentation strategy—that is, integrating human and machine work, rather than replacing humans entirely. In our survey, only 22% of executives indicated that they considered reducing head count as a primary benefit of AI.

We believe that every large company should be exploring cognitive technologies. There will be some bumps in the road, and there is no room for complacency on issues of workforce

displacement and the ethics of smart machines. But with the right planning and development, cognitive technology could usher in a golden age of productivity, work satisfaction, and prosperity.

Further Reading

- **"Big Idea: The Business of Artificial Intelligence,"** by Erik Brynjolfsson and Andrew McAfee, HBR.org/ai
- **"Inside Facebook's AI Workshop,"** by Scott Berinato, HBR.org/ai
- **"AI Can Be a Troublesome Teammate,"** by Kurt Gray, HBR.org/ai

Originally published in January–February 2018. Reprint R1801H

Why Every Organization Needs an Augmented Reality Strategy

by Michael E. Porter and James E. Heppelmann

THERE IS A FUNDAMENTAL DISCONNECT between the wealth of digital data available to us and the physical world in which we apply it. While reality is three-dimensional, the rich data we now have to inform our decisions and actions remains trapped on two-dimensional pages and screens. This gulf between the real and digital worlds limits our ability to take advantage of the torrent of information and insights produced by billions of smart, connected products (SCPs) worldwide.

Augmented reality, a set of technologies that superimposes digital data and images on the physical world, promises to close this gap and release untapped and uniquely human capabilities. Though still in its infancy, AR is poised to enter the mainstream; according to one estimate, spending on AR technology will hit $60 billion in 2020. AR will affect companies in every industry and many other types of organizations, from universities to social enterprises. In the coming months and years, it will transform how we learn, make decisions, and interact with the

physical world. It will also change how enterprises serve customers, train employees, design and create products, and manage their value chains, and, ultimately, how they compete.

In this article we describe what AR is, its evolving technology and applications, and why it is so important. Its significance will grow exponentially as SCPs proliferate, because it amplifies their power to create value and reshape competition. AR will become the new interface between humans and machines, bridging the digital and physical worlds. While challenges in deploying it remain, pioneering organizations, such as Amazon, Facebook, General Electric, Mayo Clinic, and the U.S. Navy, are already implementing AR and seeing a major impact on quality and productivity. Here we provide a road map for how companies should deploy AR and explain the critical choices they will face in integrating it into strategy and operations.

What Is Augmented Reality?

Isolated applications of AR have been around for decades, but only recently have the technologies required to unleash its potential become available. At the core, AR transforms volumes of data and analytics into images or animations that are overlaid on the real world. Today most AR applications are delivered through mobile devices, but increasingly delivery will shift to hands-free wearables such as head-mounted displays or smart glasses. Though many people are familiar with simple AR entertainment applications, such as Snapchat filters and the

game Pokémon Go, AR is being applied in far more consequential ways in both consumer and business-to- business settings. For example, AR "heads-up" displays that put navigation, collision warning, and other information directly in drivers' line of sight are now available in dozens of car models. Wearable AR devices for factory workers that superimpose production-assembly or service instructions are being piloted at thousands of companies. AR is supplementing or replacing traditional manuals and training methods at an ever-faster pace.

More broadly, AR enables a new information-delivery paradigm, which we believe will have a profound impact on how data is structured, managed, and delivered on the internet. Though the web transformed how information is collected, transmitted, and accessed, its model for data storage and delivery—pages on flat screens—has major limits: It requires people to mentally translate 2-D information for use in a 3-D world. That isn't always easy, as anyone who has used a manual to fix an office copier knows. By superimposing digital information directly on real objects or environments, AR allows people to process the physical and digital simultaneously, eliminating the need to mentally bridge the two. That improves our ability to rapidly and accurately absorb information, make decisions, and execute required tasks quickly and efficiently.

Idea in Brief

The Problem

While the physical world is three-dimensional, most data is trapped on 2-D screens and pages. This gulf between the real and digital worlds limits our ability to make the best use of the volumes of information available to us.

The Solution

Augmented reality solves this problem by superimposing digital images and data on real objects. By putting information directly into the context in which we'll apply it, AR speeds our ability to absorb and act on it.

The Outcome

Pioneering organizations, including GE, Mayo Clinic, and the U.S. Navy, are using AR to improve productivity, quality, and training. By combining the strengths of humans and machines, AR will dramatically increase value creation.

AR displays in cars are a vivid illustration of this. Until recently, drivers using GPS navigation had to look at a map on a flat screen and then figure out how to apply it in the real world. To take the correct exit from a busy rotary, for example, the driver needed to shift his or her gaze between the road and the screen and mentally connect the image on the map to the proper turnoff. AR heads-up displays lay navigational images directly over what the driver sees through the windshield. This reduces the mental effort of applying the information, prevents distraction, and minimizes driver error, freeing people to focus on the road. (For more on this, see the sidebar "Enhancing Human Decision Making.")

Enhancing Human Decision Making

At its core, the power of augmented reality grows out of the way humans process information. We access information through each of our five senses—but at different rates. Vision provides us with the most information by far: An estimated 80% to 90% of the information humans get is accessed through vision.

The ability to absorb and process information is limited by our mental capacity. The demand on this capacity is referred to as "cognitive load." Each mental task we undertake reduces the capacity available for other, simultaneous tasks.

Cognitive load depends on the mental effort required to process a given type of information. For example, reading instructions from a computer screen and acting on them creates a greater cognitive load than hearing those same instructions, because the letters must be translated into words and the words interpreted. Cognitive load also depends on "cognitive distance," or the gap between the form in which information is presented and the context in which it is applied. Consider what happens when someone refers to a smartphone for directions while driving. The driver must consume the information from the screen, retain that information in working memory, translate the directions into the physical environment in front of him, and then act on those directions, all while operating the vehicle. There is significant cognitive distance between the digital information on the screen and the physical context in which information is applied. Dealing with this distance creates cognitive load.

The combination of the speed at which information is transmitted and absorbed and the cognitive distance involved in applying it lies at the root of the much-repeated phrase "A picture is worth a thousand words." When we look at the physical world, we absorb a huge amount and variety of information almost instantaneously. By the same token, an image or picture that superimposes information on the physical world,

placing it in context for us, reduces cognitive distance and minimizes cognitive load.

This explains why AR is so powerful. There is no better graphical user interface than the physical world we see around us when it is enhanced by a digital overlay of relevant data and guidance where and when they are needed. AR eliminates dependence on out-of-context and hard-to-process 2-D information on pages and screens while greatly improving our ability to understand and apply information in the real world.

Converging physical and digital

Augmented reality reduces the mental effort needed to connect digital information about the physical world with the context it applies to.

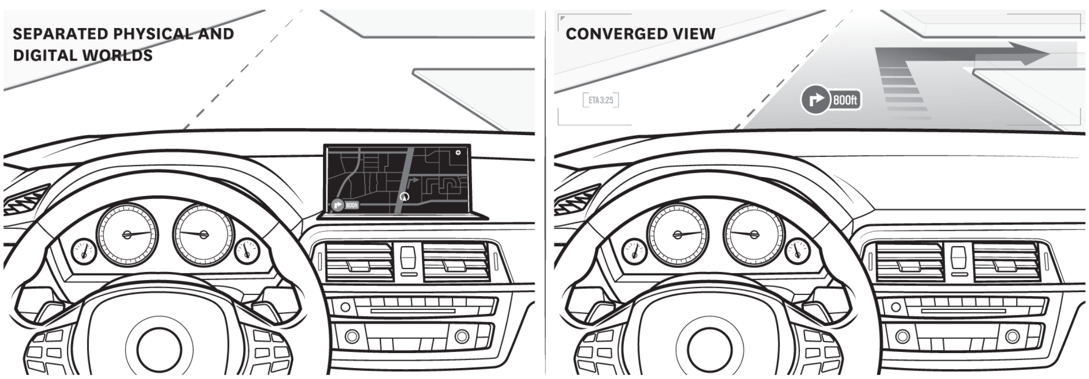

Mentally transposing GPS images onto the road ahead is demanding and prone to errors.

AR superimposes digital data directly on the real world.

AR is making advances in consumer markets, but its emerging impact on human performance is even greater in industrial settings. Consider how Newport News Shipbuilding, which designs and builds U.S. Navy aircraft carriers, uses AR near the end of its manufacturing process to inspect a ship, marking for removal steel construction structures that are not part of the

finished carrier. Historically, engineers had to constantly compare the actual ship with complex 2-D blueprints. But with AR, they can now see the final design superimposed on the ship, which reduces inspection time by 96%—from 36 hours to just 90 minutes. Overall, time savings of 25% or more are typical for manufacturing tasks using AR.

AR's Key Capabilities

As we've previously explained (see "How Smart, Connected Products Are Transforming Competition," HBR, November 2014), the SCPs spreading through our homes, workplaces, and factories allow users to monitor product operations and conditions in real time, control and customize product operations remotely, and optimize product performance using real-time data. And in some cases, intelligence and connectivity allow SCPs to be fully autonomous.

AR powerfully magnifies the value created by those capabilities. Specifically, it improves how users visualize and therefore access all the new monitoring data, how they receive and follow instructions and guidance on product operations, and even how they interact with and control the products themselves.

Visualize

AR applications provide a sort of X-ray vision, revealing internal features that would be difficult to see otherwise. At the medical

device company AccuVein, for instance, AR technology converts the heat signature of a patient's veins into an image that is superimposed on the skin, making the veins easier for clinicians to locate. This dramatically improves the success rate of blood draws and other vascular procedures. AR more than triples the likelihood of a successful needle stick on the first try and reduces the need for "escalations" (calling for assistance, for example) by 45%.

Bosch Rexroth, a global provider of power units and controls used in manufacturing, uses an AR-enhanced visualization to demonstrate the design and capabilities of its smart, connected CytroPac hydraulic power unit. The AR application allows customers to see 3-D representations of the unit's internal pump and cooling options in multiple configurations and how subsystems fit together.

Instruct and guide

AR is already redefining instruction, training, and coaching. These critical functions, which improve workforce productivity, are inherently costly and labor-intensive and often deliver uneven results. Written instructions for assembly tasks, for instance, are frequently hard and time-consuming to follow. Standard instructional videos aren't interactive and can't adapt to individual learning needs. In-person training is expensive and requires students and teachers to meet at a common site, sometimes repeatedly. And if the equipment about which

students are being taught isn't available, they may need extra training to transfer what they've learned to a real-world context.

AR addresses those issues by providing real-time, on-site, step-by-step visual guidance on tasks such as product assembly, machine operation, and warehouse picking. Complicated 2-D schematic representations of a procedure in a manual, for example, become interactive 3-D holograms that walk the user through the necessary processes. Little is left to the imagination or interpretation.

At Boeing, AR training has had a dramatic impact on the productivity and quality of complex aircraft manufacturing procedures. In one Boeing study, AR was used to guide trainees through the 50 steps required to assemble an aircraft wing section involving 30 parts. With the help of AR, trainees completed the work in 35% less time than trainees using traditional 2-D drawings and documentation. And the number of trainees with little or no experience who could perform the operation correctly the first time increased by 90%.

AR-enabled devices can also transmit what an on-site user is seeing to a remote expert, who can respond with immediate guidance. In effect, this instantly puts the expert at the user's side, regardless of location. This capability not only improves worker performance but substantially reduces costs—as Lee Company, which sells and services building systems, has discovered. It uses AR to help its field technicians with installations and repairs. A remote expert can see what the tech

is viewing through his or her AR device, guide the tech through the work to be done, and even annotate the tech's view with instructions. Getting expert support from a central location in real time has increased Lee's tech utilization dramatically. And, by reducing the number of repeat visits, Lee saves more than $500 per technician per month in labor and travel costs. The company calculates a return of $20 on every dollar invested in AR.

Interact

Traditionally, people have used physical controls such as buttons, knobs, and, more recently, built-in touchscreens to interact with products. With the rise of SCPs, apps on mobile devices have increasingly replaced physical controls and allowed users to operate products remotely.

AR takes the user interface to a whole new level. A virtual control panel can be superimposed directly on the product and operated using an AR headset, hand gestures, and voice commands. Soon, users wearing smart glasses will be able to simply gaze at or point to a product to activate a virtual user interface and operate it. A worker wearing smart glasses, for instance, will be able to walk a line of factory machines, see their performance parameters, and adjust each machine without physically touching it.

The interact capability of AR is still nascent in commercial products but is revolutionary. Reality Editor, an AR app developed by the Fluid Interfaces group at MIT's Media Lab,

provides a glimpse of how it is rapidly evolving. Reality Editor makes it easy to add an interactive AR experience to any SCP. With it, people can point a smartphone or a tablet at an SCP (or, eventually, look at it through smart glasses), “see” its digital interfaces and the capabilities that can be programmed, and link those capabilities to hand gestures or voice commands or even to another smart product. For example, Reality Editor can allow a user to see a smart light bulb’s controls for color and intensity and set up voice commands like “bright” and “mood” to activate them. Or different settings of the bulb can be linked to buttons on a smart light switch the user can place anywhere that’s convenient.

The technologies underpinning these capabilities are still emerging, but the accuracy of voice commands in noisy environments is improving, and advances in gesture and gaze tracking have been rapid. GE has already tested the use of voice commands in AR experiences that enable factory workers to perform complex wiring processes in wind turbines—and has achieved a 34% increase in productivity.

Combining AR and Virtual Reality

AR’s well-known cousin, virtual reality, is a complementary but distinct technology. While AR superimposes digital information on the physical world, VR replaces physical reality with a computer- generated environment. Though VR is used mostly for entertainment applications, it can also replicate physical

settings for training purposes. It is especially useful when the settings involved are hazardous or remote. Or, if the machinery required for training is not available, VR can immerse technicians in a virtual environment using holograms of the equipment. So when needed, VR adds a fourth capability—simulate—to AR's core capabilities of visualize, instruct, and interact.

AR will be far more widely applied in business than VR will. But in some circumstances, combining AR and VR will allow users to transcend distance (by simulating faraway locations), transcend time (by reproducing historical contexts or simulating possible future situations), and transcend scale (by allowing users to engage with environments that are either too small or too big to experience directly). What's more, bringing people together in shared virtual environments can enhance comprehension, teamwork, communication, and decision making.

Ford, for example, is using VR to create a virtual workshop where geographically dispersed engineers can collaborate in real time on holograms of vehicle prototypes. Participants can walk around and go inside these life-size 3-D holograms, working out how to refine design details such as the position of the steering wheel, the angle of the dashboard, and the location of instruments and controls without having to build an expensive physical prototype and get everyone to one location to examine it.

The U.S. Department of Homeland Security is going a step further by combining AR instructions with VR simulations to train personnel in responding to emergency situations such as explosions. This reduces costs and—in cases in which training in real environments would be dangerous—risk. The energy multinational BP overlays AR training procedures on VR simulations that replicate specific drilling conditions, like temperature, pressure, topography, and ocean currents, and that instruct teams on operations and help them practice coordinated emergency responses to disasters without high costs or risk.

How AR Creates Value

AR creates business value in two broad ways: first, by becoming part of products themselves, and second, by improving performance across the value chain—in product development, manufacturing, marketing, service, and numerous other areas.

AR as a product feature

The capabilities of AR play into the growing design focus on creating better user interfaces and ergonomics. The way products convey important operational and safety information to users has increasingly become a point of differentiation (consider how mobile apps have supplemented or replaced embedded screens in products like Sonos audio players). AR is poised to rapidly improve such interfaces.

Dedicated AR heads-up displays, which have only recently been incorporated into automobiles, have been a key feature in elite military products, such as fighter jets, for years and have been adopted in commercial aircraft as well. These types of displays are too expensive and bulky to integrate into most products, but wearables such as smart glasses are a breakthrough interface with wide-ranging implications for all manufacturers. With smart glasses, a user can see an AR display on any product enabled to communicate with them.

If you view a kitchen oven through smart glasses, for example, you might see a virtual display that shows the baking temperature, the minutes remaining on the timer, and the recipe you are following. If you approach your car, an AR display might show you that it is locked, that the fuel tank is nearly full, and that the left-rear tire's pressure is low.

Because an AR user interface is purely software based and delivered via the cloud, it can be personalized and can continually evolve. The incremental cost of providing such an interface is low, and manufacturers also stand to save considerable amounts when traditional buttons, switches, and dials are removed. Every product manufacturer needs to carefully consider the disruptive impact that this next-generation interface may have on its offering and competitive positioning.

AR and the value chain

The effects of AR can already be seen across the value chain, but they are more advanced in some areas than in others. In general, visualize and instruct/guide applications are now having the greatest impact on companies' operations, while the interact capability is still emerging and in pilot testing.

Product development. Though engineers have been using computer-aided design (CAD) capabilities to create 3-D models for 30 years, they have been limited to interacting with those models through 2-D windows on their computer screens, which makes it harder for them to fully conceptualize designs. AR allows 3-D models to be superimposed on the physical world as holograms, enhancing engineers' ability to evaluate and improve designs. For example, a life-size 3-D hologram of a construction machine can be positioned on the ground, and engineers can walk around it, peer under and over it, and even go inside it to fully appreciate the sight lines and ergonomics of its design at full scale in its intended setting.

AR also lets engineers superimpose CAD models on physical prototypes to compare how well they match. Volkswagen is using this technique—which makes any difference between the latest design and the prototype visually obvious—to check alignment in digital design reviews. This improves the accuracy of the quality assurance process, in which engineers previously had to painstakingly compare 2-D drawings with prototypes, and makes it five to 10 times faster.

We expect that in the near future AR-enabled devices such as phones and smart glasses, with their embedded cameras, accelerometers, GPS, and other sensors, will increasingly inform product design by exposing when, where, and how users actually interact with the product—how often a certain repair sequence is initiated, for example. In this way the AR interface will become an important source of data.

Manufacturing. In manufacturing, processes are often complex, requiring hundreds or even thousands of steps, and mistakes are costly. As we've learned, AR can deliver just the right information the moment it's needed to factory workers on assembly lines, reducing errors, enhancing efficiency, and improving productivity.

In factories, AR can also capture information from automation and control systems, secondary sensors, and asset management systems and make visible important monitoring and diagnostic data about each machine or process. Seeing information such as efficiency and defect rates in context helps maintenance technicians understand problems and prompts factory workers to do proactive maintenance that may prevent costly downtime.

Iconics, which specializes in automation software for factories and buildings, has begun to integrate AR into its products' user interfaces. By attaching relevant information to the physical location where it will be best observed and understood, the AR interfaces enable more-efficient monitoring of machines and processes.

Logistics. Warehouse operations are estimated to account for about 20% of all logistics costs, while picking items from shelves represents up to 65% of warehouse costs. In most warehouses, workers still perform this task by consulting a paper list of things to collect and then searching for them. This method is slow and error-prone.

The logistics giant DHL and a growing number of other companies are using AR to enhance the efficiency and accuracy of the picking process. AR instructions direct workers to the location of each product to be pulled and then suggest the best route to the next product. At DHL this approach has led to fewer errors, more-engaged workers, and productivity gains of 25%. The company is now rolling out AR-guided picking globally and testing how AR can enhance other types of warehouse operations, such as optimizing the position of goods and machines in layouts. Intel is also using AR in warehouses and has achieved a 29% reduction in picking time, with error rates falling to near zero. And the AR application is allowing new Intel workers to immediately achieve picking speeds 15% faster than those of workers who've had only traditional training.

Marketing and sales. AR is redefining the concept of showrooms and product demonstrations and transforming the customer experience. When customers can see virtually how products will look or function in a real setting before buying them, they have more-accurate expectations, more confidence about their purchase decisions, and greater product satisfaction. Down the

road, AR may even reduce the need for brick-and-mortar stores and showrooms altogether.

When products can be configured with different features and options—which can make them difficult and costly to stock—AR is a particularly valuable marketing tool. The construction products company AZEK, for instance, uses AR to show contractors and consumers how its decking and paver products look in various colors and arrangements. Customers can also see the simulations in context: If you look at a house through a phone or a tablet, the AR app can add a deck onto it. The experience reduces any uncertainty customers might feel about their choices and shortens the sales cycle.

In e-commerce, AR applications are allowing online shoppers to download holograms of products. Wayfair and IKEA both offer libraries with thousands of 3-D product images and apps that integrate them into a view of an actual room, enabling customers to see how furniture and decor will look in their homes. IKEA also uses its app to collect important data about product preferences in different regions.

After-sales service. This is a function where AR shows huge potential to unlock the value-creating capabilities of SCPs. AR assists technicians serving customers in the field in much the same way it helps workers in factories: by showing predictive analytics data generated by the product, visually guiding them through repairs in real time, and connecting them with remote experts who can help optimize procedures. For example, an AR

dashboard might reveal to a field technician that a specific machine part will most likely fail within a month, allowing the tech to preempt a problem for the customer by replacing it now.

At KPN, a European telecommunications service provider, field engineers conducting remote or on-site repairs use AR smart glasses to see a product's service-history data, diagnostics, and location-based information dashboards. These AR displays help them make better decisions about how to resolve issues, producing an 11% reduction in overall costs for service teams, a 17% decrease in work-error rates, and higher repair quality.

Xerox used AR to connect field engineers with experts instead of providing service manuals and telephone support. First-time fix rates increased by 67%, and the engineers' efficiency jumped by 20%. Meanwhile, the average time it took to resolve problems dropped by two hours, so staffing needs fell. Now Xerox is using AR to connect remote technical experts directly with customers. This has increased by 76% the rate at which technical problems are resolved by customers without any on-site help, cutting travel costs for Xerox and minimizing downtime for customers. Perhaps not surprisingly, Xerox has seen its customer satisfaction rates rise to 95%.

Human resources. Early AR adopters like DHL, the U.S. Navy, and Boeing have already discovered the power of delivering step-by-step visual worker training on demand through AR. AR allows instruction to be tailored to a particular worker's experience or to reflect the prevalence of particular errors. For example, if

someone repeatedly makes the same kind of mistake, he can be required to use AR support until his work quality improves. At some companies, AR has reduced the training time for new employees in certain kinds of work to nearly zero and lowered the skill requirements for new hires.

This is especially advantageous for the package delivery company DHL, which faces surges in demand during peak seasons and is heavily dependent on the effective hiring and training of temporary workers. By providing real-time training and hands-on guidance on navigating warehouses and properly packing and sorting materials, AR has reduced DHL's need for traditional instructors and increased the onboarding speed for new employees.

AR and Strategy

AR will have a widespread impact on how companies compete. As we've explained in our previous HBR articles, SCPs are changing the structure of almost all industries as well as the nature of competition within them—often expanding industry boundaries in the process. SCPs give rise to new strategic choices for manufacturers, ranging from what functionality to pursue and how to manage data rights and security, to whether to expand a company's scope of products and compete in smart systems.

The increasing penetration of AR, along with its power as the human interface with SCP technologies, raises some new

strategic questions. While the answers will reflect each company's business and unique circumstances, AR will become more and more integral to every firm's strategy.

Here are the essential questions companies face:

1. What is the range of AR opportunities in the industry, and in what sequence should they be pursued? Companies must weigh AR's potential impact on customers, product capabilities, and the value chain.

2. How will AR reinforce a company's product differentiation? AR opens up multiple differentiation paths. It can create companion experiences that expand the capabilities of products, give customers more information, and increase product loyalty. AR interfaces that enhance products' functionality or ease of use can be big differentiators, as can those that substantially improve product support, service, and uptime. And AR's capacity to provide new kinds of feedback on how customers use products can help companies uncover further opportunities for product differentiation.

The right differentiation path will depend on a company's existing strategy; what competitors are doing; and the pace of technology advances, especially in hardware.

3. Where will AR have the greatest impact on cost reduction? AR enables new efficiencies that every firm must explore. As we've noted, it can significantly lower the cost of training, service, assembly, design, and other parts of the value chain. It can also

substantially cut manufacturing costs by reducing the need for physical interfaces.

Each company will need to prioritize AR-driven cost-reduction efforts in a way that's consistent with its strategic positioning. Firms with sophisticated products will need to capitalize on AR's superior and low-cost interface, while many commodity producers will focus on operational efficiencies across the value chain. In consumer industries and retail, marketing-related visualize applications are the most likely starting point. In manufacturing, instruct applications are achieving the most immediate payoff by addressing inefficiencies in engineering, production, and service. And AR's interact capability, though still emerging, will be important across all industries with products that have customization and complex control capabilities.

4. Should the company make AR design and deployment a core strength, or will outsourcing or partnering be sufficient? Many firms are scrambling to access the digital talent needed for AR development, which is in short supply. One skill in great demand is user experience or user interface (UX/UI) design. It's critical to present 3-D digital information in ways that make it easy to absorb and act on; companies want to avoid making a stunning but unhelpful AR experience that defeats its core purpose. Effective AR experiences also require the right content, so people who know how to create and manage it—another novel skill—are crucial too. Digital modeling capabilities and knowledge of how to apply them in AR applications are key as well.

Over time we expect companies to create teams dedicated to AR, just as they set up such teams to build and run websites in the 1990s and 2000s. Dedicated teams will be needed to establish the infrastructure that will allow this new medium to flourish and to develop and maintain the AR content. Many firms have started to build AR skills in-house, but few have mastered them yet.

Whether to hire and train AR employees or partner with specialty software and services companies is an open question for many. Some companies have no choice but to treat AR talent as a strategic asset and invest in acquiring and developing it, given AR's potentially large impact on competition in their business. However, if AR is important but not essential to competitive advantage, firms can partner with specialty software and services companies to leverage outside talent and technology.

The challenges, time, and cost involved in building the full set of AR technologies we have described are significant, and specialization always emerges in each component. In the early stages of AR, the number of technology and service suppliers has been limited, and companies have built internal capabilities. However, best-of-breed AR vendors with turnkey solutions are starting to appear, and it will become increasingly difficult for in-house efforts to keep up with them.

5. How will AR change communications with stakeholders? AR complements existing print and 2-D digital communication

approaches and in some cases can replace them altogether. Yet we see AR as much more than just another communication channel. It is a fundamentally new means of engaging with people. Just consider the novel way it helps people absorb and act on information and instructions.

The web, which began as a way to share technical reports, ultimately transformed business, education, and social interaction. We expect that AR will do the same thing for communication—changing it in ways far beyond what we can envision today. Companies will need to think creatively about how they can use this nascent channel.

Deploying AR

AR applications are already being piloted and deployed in products and across the value chain, and their number and breadth will only grow.

Every company needs an implementation road map that lays out how the organization will start to capture the benefits of AR in its business while building the capabilities needed to expand its use. When determining the sequence and pace of adoption, companies must consider both the technical challenges and the organizational skills involved, which vary from context to context. Specifically, organizations need to address five key questions:

1. Which development capabilities will be required? Some AR experiences involve more complexity than others. Experiences that allow people to visualize products in different configurations or settings—like those created by IKEA, Wayfair, and AZEK—are a relatively easy place for companies to start. Consumers just need to be encouraged to download and launch AR apps, and only a mobile device is needed to use them.

Instruction applications, like the ones Boeing and GE employ in manufacturing, are more difficult to build and use. They require the capacity to develop and maintain dynamic 3-D digital content and often benefit greatly from the use of head-mounted displays or smart glasses, which are still in the early stages of development.

Apps that produce interactive experiences, which create significant value for both consumers and businesses, are the most challenging to develop. They also involve less-mature technology, such as voice or gesture recognition, and the need to integrate with software that controls SCPs. Most companies will start with static visualizations of 3-D models, but they should build the capability to move quickly into dynamic instructional experiences that have greater strategic impact.

2. How should organizations create digital content? Every AR experience, from the least to the most sophisticated, requires content. In some cases it's possible to repurpose existing digital content, such as product designs. Over time, however, more-

complex, dynamic contextual experiences must be built from scratch, which requires specialized expertise.

Simple applications, such as an AR-enhanced furniture catalog, may need only basic product representations. More-sophisticated business instruction applications, however, such as those used for machine repair, will require accurate and highly detailed digital product representations. Companies can create these by adapting CAD models used in product development or by using digitization techniques such as 3-D scanning. The most sophisticated AR experiences also need to tap real-time data streams from enterprise business systems, SCPs, or external data sources and integrate them into the content. To prepare for broadening the AR portfolio, companies should take an inventory of existing 3-D digital assets in CAD and elsewhere and invest in digital modeling capabilities.

3. How will AR applications recognize the physical environment? To accurately superimpose digital information on the physical world, AR technologies must recognize what they're looking at. The simplest approach is to determine the location of the AR device using, say, GPS and show relevant information for that location without anchoring it to a specific object. This is known as an "unregistered" AR experience. Vehicle heads-up navigation displays typically work this way.

Higher-value "registered" experiences anchor information to specific objects. They can do this through markers, such as bar codes, logos, or labels, which are placed on the objects and

scanned by the user with an AR device. A more powerful approach, however, uses technology that recognizes objects by comparing their shape to a catalog of 3-D models. This allows a maintenance technician, for example, to instantly recognize and interact with any type of equipment he or she is responsible for maintaining and to do so from any angle. While markers are a good starting point, shape-recognition technologies are advancing quickly, and organizations will need the capability to use them to tap into many of the highest-value AR applications.

4. **What AR hardware is required?** AR experiences aimed at broad consumer audiences have typically been designed for smartphones, taking advantage of their simplicity and ubiquity. For more- sophisticated experiences, companies use tablets, which offer larger screens, better graphics, and greater processing power. Since tablet penetration is lower, companies will often provide them to users. For certain high-value applications—notably those in aircraft and automobiles—manufacturers are building dedicated AR heads-up displays into their products—a costly approach.

Eventually, however, most AR applications for service, manufacturing, and even product interfaces will require head-mounted displays that free users' hands. This technology is currently both immature and expensive, but we expect that affordable smart glasses will become widely available in the next few years and will play a major part in releasing AR's full power. Microsoft, Google, and Apple now offer AR technologies

optimized for their own devices. However, most organizations should take a cross-platform approach that allows AR experiences to be deployed across multiple brands of phones and tablets and should make sure they're ready for smart glasses when they arrive.

5. Should you use a software-development or a content-publishing model?

Many early AR experiences have been delivered through stand-alone software applications that are downloaded, complete with digital content, to a phone or a tablet. This approach creates reliable, high-resolution experiences and allows organizations to make apps that don't require internet connectivity. The problem with this model is that any change to the AR experience requires software developers to rewrite the app, which can create expensive bottlenecks.

An emerging alternative uses commercial AR-publishing software to create AR content and host it in the cloud. The AR experience can then be downloaded on demand using a general-purpose app running on an AR device. Like website content, the AR content can be updated or supplemented without changing the software itself—an important benefit when large amounts of information and frequent content changes are involved. The content- publishing model will become common as more and more machines and products include real-time AR interaction and control. A content-publishing capability is essential to scaling AR up across the organization.

The Broader Impact

The digital revolution, with its SCPs and explosion of data, is unleashing productivity and unlocking value across the economy. Increasingly, the constraint is not a lack of data and knowledge but how to assimilate and act on them—in other words, the interface with humans. AR is emerging as a leading solution to this challenge.

At the same time, the rapid evolution of machine learning and automation is raising serious concerns about human opportunity. Will there be enough jobs for everyone, especially for people without advanced education and knowledge? In a world of artificial intelligence and robots, will humans become obsolete?

It is easy to conclude that new technology diminishes human opportunity. Yet new inventions have been replacing human labor for centuries, and they have led to growth in employment, not a decline. Technology has dramatically increased our productivity and our standard of living. It has given rise to new kinds of offerings that meet new needs and require new types of workers. Many of today's jobs involve products and services that did not even exist a hundred years ago. A lesson of history is that today's digital revolution will generate new waves of innovation and new kinds of work that we cannot yet imagine.

The role of humans in this future is misunderstood. People have unique strengths that machines and algorithms will not replicate anytime soon. We have sophisticated motor skills—well beyond what robots are capable of today—that allow us to do the

subtle manipulation that's needed in, say, replacing a machine part or wiring a turbine. Even relatively less skilled work, such as drawing blood, pruning a garden, or repairing a flat tire, requires human dexterity and defies automation. Human cognition adapts instantaneously to novel situations; people easily adjust the way they interpret information, solve problems, exercise judgment, and take action to suit their circumstances. Humans have flexibility, imagination, intuition, and creative ability that for the foreseeable future are beyond the reach of any machine.

How Does Augmented Reality Work?

Augmented reality starts with a camera-equipped device—such as a smartphone, a tablet, or smart glasses—loaded with AR software. When a user points the device and looks at an object, the software recognizes it through computer vision technology, which analyzes the video stream.

The device then downloads information about the object from the cloud, in much the same way that a web browser loads a page via a URL. A fundamental difference is that the AR information is presented in a 3-D "experience" superimposed on the object rather than in a 2-D page on a screen. What the user sees, then, is part real and part digital.

AR can provide a view of the real-time data flowing from products and allow users to control them by touchscreen, voice, or gesture. For example, a user might touch a stop button on the digital graphic overlay within an AR experience—or simply say the word "stop"—to send a command via the cloud to a product. An operator using an AR headset to interact with an industrial robot might see superimposed data about the robot's performance and gain access to its controls.

As the user moves, the size and orientation of the AR display automatically adjust to the shifting context. New graphical or text information comes into view while other information passes out of view. In industrial settings, users in different roles, such as a machine operator and a maintenance technician, can look at the same object but be presented with different AR experiences that are tailored to their needs.

A 3-D digital model that resides in the cloud—the object's "digital twin"—serves as the bridge between the smart object and the AR. This model is created either by using computer-aided design, usually during product development, or by using technology that digitizes physical objects. The twin then collects information from the product, business systems, and external sources to reflect the product's current reality. It is the vehicle through which the AR software accurately places and scales up-to-date information on the object.

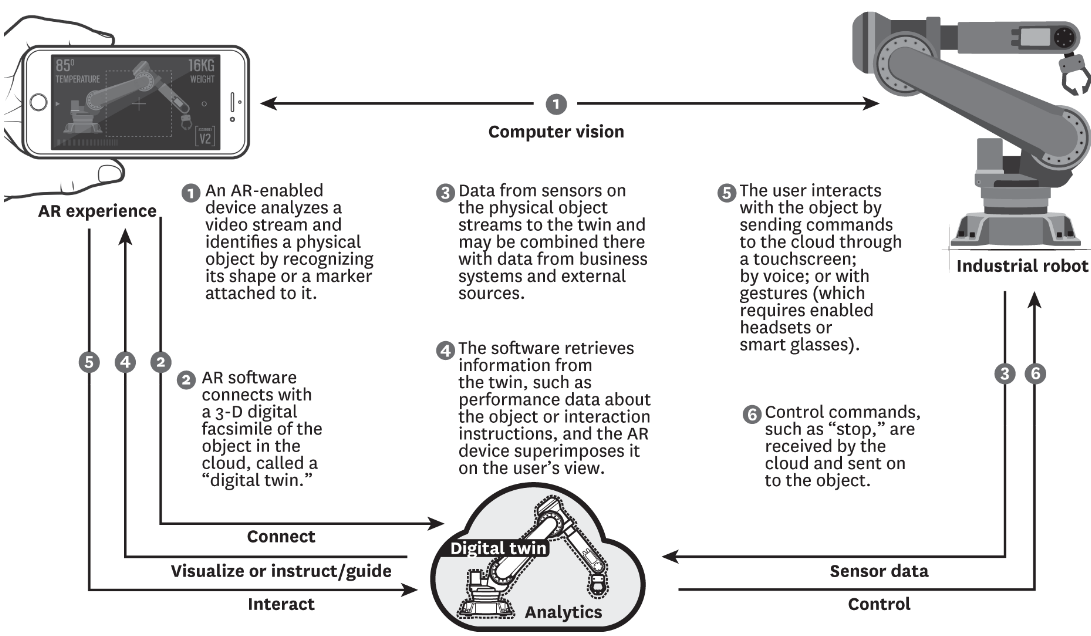

Clint Ford

While the advances in artificial intelligence and robotics are impressive, we believe that combining the capabilities of machines with humans' distinctive strengths will lead to far greater productivity and more value creation than either could generate alone. What's needed to realize this opportunity is a powerful human interface that bridges the gap between the digital and physical worlds. We see AR as a historic innovation that provides this. It helps humans enhance their own capabilities by taking full advantage of new digital knowledge and machine capabilities. It will profoundly change training and skill development, allowing people to perform sophisticated work without protracted and expensive conventional instruction — a model that is inaccessible to so many today. AR, then, enables people to better tap into the digital revolution and all it has to offer.

Originally published in November–December 2017. Reprint R1706B

Thriving in the Gig Economy

by Gianpiero Petriglieri, Susan Ashford, and Amy Wrzesniewski

HAVE YOU EVER BEEN ON a trapeze?" That's how Martha, an independent consultant, responded when we asked her to describe her work in the five years since she'd left a global consulting firm to set out on her own. She had recently tried the art, which she saw as a good metaphor for her life: the void she felt when between assignments; the exhilaration of landing the next engagement; the discipline, concentration, and grace that mastering her profession required. Trapeze artists seem to take huge risks, she explained, but a safety system—including nets, equipment, and fellow performers—supports them: "They appear to be on their own, but they're not."

Martha (whose name, like others in this article, has been changed) is part of a burgeoning segment of the workforce loosely known as the gig economy. Approximately 150 million workers in North America and Western Europe have left the relatively stable confines of organizational life—sometimes by

choice, sometimes not—to work as independent contractors. Some of this growth reflects the emergence of ride-hailing and task-oriented service platforms, but a recent report by McKinsey found that knowledge-intensive industries and creative occupations are the largest and fastest-growing segments of the freelance economy.

To learn what it takes to be successful in independent work, we recently completed an in-depth study of 65 gig workers. We found remarkably similar sentiments across generations and occupations: All those we studied acknowledged that they felt a host of personal, social, and economic anxieties without the cover and support of a traditional employer—but they also claimed that their independence was a choice and that they would not give up the benefits that came with it. Although they worried about unpredictable schedules and finances, they also felt they had mustered more courage and were leading richer lives than their corporate counterparts.

We discovered that the most effective independent workers navigate this tension with common strategies. They cultivate four types of connections—to *place, routines, purpose,* and *people*—that help them endure the emotional ups and downs of their work and gain energy and inspiration from their freedom. As the gig economy grows worldwide, these strategies are increasingly relevant. Indeed, we believe they may also be helpful to any corporate employees who are working more

autonomously, from home or a remote office, or who feel they might one day want—or need—to jump into a freelance career.

Produce or Perish

The first thing we realized when we began interviewing independent consultants and artists was that the stakes of independent work are enormously high—not just financially but also existentially. Unshackled from managers and corporate norms, people can choose assignments that make the most of their talents and reflect their true interests. They feel ownership over what they produce and over their entire professional lives. One study participant told us, "I can be the most I've ever been myself in any job."

However, the price of such freedom is a precariousness that seems not to subside over time. Even the most successful, well-established people we interviewed still worry about money and reputation and sometimes feel that their identity is at stake. You can't keep calling yourself a consultant, for example, if clients stop asking for your services. A well-published writer told us, "You become your work. If you write a good book . . . it's really great, and when you don't achieve it, you have to accept . . . that failure might define who you are to yourself." An artist agreed: "There's no arriving. That's a myth."

Idea in Brief

Approximately 150 million people in North America and Western Europe now work as independent contractors, most of them in knowledge-intensive industries and creative occupations. The authors studied 65 of them in depth and learned that although they feel a host of personal, social, and economic anxieties without the cover and support of a traditional employer, they also say they chose independence and wouldn't give up the benefits that come with it.

Many of these workers have created a "holding environment" for themselves by establishing four connections: (1) place, in the form of idiosyncratic, dedicated workspaces that allow easy access to the tools of their owners' trades; (2) routines that streamline workflow and incorporate personal care; (3) purpose, to create a bridge between personal interests and motivations and a need in the world; and (4) people to whom they turn for reassurance and encouragement. These connections help independent workers sustain productivity, endure their anxieties, and even turn those feelings into sources of creativity and growth.

For this reason, productivity is an intense preoccupation for everyone we interviewed. It provides self-expression and an antidote to precariousness. Interestingly, however, the people we talked with aren't just focusing on getting things done and sold. They care about both being *at work*—having the discipline to regularly generate products or services that find a market—and being *into their work:* having the courage to stay fully invested in the process and output of that labor.

Sustaining productivity is a constant struggle. Distress and distractions can erode it, and both impediments abound in people's working lives. One executive coach gave a poignant description of an unproductive day: "It's when there is so much

to do that I’m disorganized and can’t get my act together. [In the evening,] the same e-mails I opened in the morning are still open. The documents I wanted to get done are not done. I got distracted and feel like I wasted time.” A day like that, he said, leaves him full of self-doubt.

When we asked interviewees the secret to getting through such days and ultimately sustaining productivity as they defined it, we discovered a paradox at the heart of their answers. They all want to preserve their independence and, in many cases, even their unsettledness (which one consultant described as the key to continued learning and “keeping my edge”), but they also spend a great deal of time developing a “holding environment”—a physical, social, and psychological space for their work.

This concept—first used by the British psychoanalyst Donald Winnicott to describe how attentive caregivers facilitate children’s development by buffering them against distress and creating room for experimentation—has since been employed in the field of adult development to refer to conditions in which people can be their best and grow. Corporate employees, of course, can find them with a good boss in a solid organization. But for independent workers, a holding environment is less a gift than an accomplishment; it must be cultivated, and it can be lost.

So they create these environments for themselves by establishing and maintaining what we call “liberating

connections"—because they both *free* people up to be individually creative and *bind* them to work so that their output doesn't wane.

The Four Connections

Place

Disconnected from a corporate office, the people we interviewed find places to work that protect them from outside distractions and pressures and help them avoid feeling rootless. Though many claimed their work was portable, they all still seemed to have somewhere to retreat. One writer told us, "People fail because they don't create a space and time to do whatever it is they need to do."

We visited many of these spaces in person and noticed several similarities among them. They feel confined—almost uncomfortably so in the case of some artists. They are used consistently for all substantive work. They allow easy access to the tools of the owner's trade and to little else. And they're dedicated to work; people usually leave them once their daily tasks are done. One software engineer, whose home office has all these features, described it as a "fighter pilot cockpit," where everything he needs is within arm's reach. "Sometimes it's claustrophobic," he explained, but "when I'm there, the open space is in my mind."

Despite these commonalities, each workspace is also unique, with a location, furniture, supplies, and decorations that reflect the idiosyncrasy of its owner's work. These places are not just protective cocoons for the working self—they evoke it, too. Karla, an independent consultant who initially told us she could do work "wherever I show up and am doing something that has positive impact in the world," eventually admitted that her home office is where she goes to avoid distraction and find inspiration, literally surrounded by her current and potential projects, arranged in visible and accessible piles. "When I walk through that door, I step into a space that embraces all the different aspects of myself," she told us. "I feel at home in there." Without that place and the space it gives her, Karla explained, she would probably be too sensitive to external demands and thus less focused and free.

Routines

In organizations, routines are often associated with safety or boring bureaucracy. However, a growing body of research has shown that elite athletes, scientific geniuses, popular artists, and even everyday workers use routines to enhance focus and performance. The professionals we spoke with tend to rely on them in the same way.

Some routines improve people's workflow: keeping a schedule; following a to-do list; beginning the day with the most challenging work or with a client call; leaving a sentence

incomplete in an unfinished manuscript to make an easy start the next day; sweeping the studio floor while reflecting on a new piece. Other routines, usually involving sleep, meditation, nutrition, or exercise, incorporate personal care into people's working lives. Both kinds often have a ritual element that enhances people's sense of order and control in uncertain circumstances.

One consultant we interviewed takes a bath every morning and visualizes what she wants to accomplish while she soaks. Another consultant, Matthew, who specializes in helping boards focus on innovation, keeps a strict daily schedule: "I'm up at 6:00 and there's exercise. I pack my wife's lunch. We pray. She's out the door around 8:00. I'm in my office by 8:30, and I do work where there's deeper thought required—design or writing—in the morning. That's when I'm at my best. Then in the afternoon I schedule phone calls, more of the business or financial things that need to be done." This discipline even extends to his wardrobe: "I always get dressed for the office. Most days in summer I wear shorts when I'm not on the road, but still I shower and shave as if I were going to a workplace separate from home."

That may sound rigid, but it helps Matthew pour himself into his work. He and other successful independent workers seem to follow the advice of the French novelist Gustave Flaubert: "Be regular and orderly in your life . . . so that you may be violent and original in your work."

Purpose

For most people in our study, striking out on their own initially involved doing whatever work would allow them to find a footing in the market. But they were adamant that succeeding means taking only work that clearly connects to a broader purpose. All could articulate why their work, or at least their best work—be it to empower women through film, expose harmful marketing practices, sustain the American folk music tradition, or help corporate leaders succeed with integrity— is more than a means of earning a living. Purpose creates a bridge between their personal interests and motivations and a need in the world. Matthew, for example, said that although at first he felt "a certain desperation around having clients and making an income," over time his view of success shifted "to one that is a lot about living a life of service to others and making the planet a better place."

An executive coach we interviewed told us that purpose keeps her steady, inspired, and inspiring. "A big distinction between successful independents and the ones who aren't or go back [to corporate jobs] is getting to that place of knowing what you're meant to do. That gives me resilience for the ups and downs. It gives me the strength to decline work that isn't in alignment. It gives me a quality of authenticity and confidence that clients are drawn to. It's helpful to building or maintaining the business and serving the people I am here to serve."

We found that purpose, like the other connections, both binds and frees people by orienting and elevating their work.

People

Humans are social creatures. Studies in corporate settings have long demonstrated how important other people are to our careers—as role models who show us who we might become, and as peers who help us progress by sharing our path. Researchers have also warned about a “loneliness epidemic” hitting the workplace, for which independent workers can certainly be at even greater risk.

But those we interviewed are keenly aware of the dangers of social isolation and strive to avoid it. Though many are ambivalent about formal peer groups, which they often see as insipid substitutes for collegiality, all reported having people they turn to for reassurance and encouragement. Sometimes these are direct role models or supportive collaborators; in other cases they’re family members, friends, or contacts in similar fields, who can’t always offer specific work advice but nevertheless help our study participants push through challenging times and embolden them to take the risks their work entails.

Matthew, for example, noted that reaching out to people in his inner circle helps calm his anxiety: “If I were just left on my own, I could sit here in the office and go down a rat hole. You’re left to your own inner voice, and it spirals down into

ruminating." Karla told us that she, too, regularly turns to a handful of peers with whom she's close. "All the work I do in the independent economy comes through these connections," she said. But their help goes well beyond referrals. "My ability to process, develop, and grow as a human being and understand who I am in the work I'm doing comes from the conversations that I have with these folks," she explained. "These people are how I know what I'm supposed to be doing."

Redefining Success

In popular management tales, career success usually comes with security and equanimity. For independent workers, however, both are ultimately elusive. And yet most of those we studied told us they feel successful.

Our conclusion is that people in the gig economy must pursue a different kind of success—one that comes from finding a balance between predictability and possibility, between viability (the promise of continued work) and vitality (feeling present, authentic, and alive in one's work). Those we interviewed do so by building holding environments around place, routines, purpose, and people, which help them sustain productivity, endure their anxieties, and even turn those feelings into sources of creativity and growth. "There's a sense of confidence that comes from a career as a self-employed person," one consultant told us. "You can feel that no matter

how bad it gets, I can overcome this. I can change it. I can operate more from a place of choice as opposed to a place of need."

Many we spoke to believe they wouldn't be able to find the same mental space or strength in a traditional workplace. Martha, the consultant who compared herself to a trapeze artist, recalled that she became "much more successful professionally" and "much more comfortable in my identity personally" when a trusted counselor helped her reframe—and own—her struggle, rather than seek ways to evade it. "She helped me understand that I could think of myself, which I now do, as a pioneer. I don't fit in any categories that exist in organizations, and it's more effective for me to be independent." Seen this way, discomfort and uncertainty were not just tolerable but affirming—signs that she was just where she needed to be.

When we spoke, she portrayed employment as no longer an anchor she missed but a shackle she'd been fortunate enough to break. "I don't know that I would frame [my new life] as precariousness anymore," she concluded. "I would frame it as really living."

Originally published in March–April 2018. Reprint R1802M

Managing Our Hub Economy

by Marco Iansiti and Karim R. Lakhani

THE GLOBAL ECONOMY IS COALESCING around a few digital superpowers. We see unmistakable evidence that a winner-take-all world is emerging in which a small number of "hub firms"—including Alibaba, Alphabet/Google, Amazon, Apple, Baidu, Facebook, Microsoft, and Tencent—occupy central positions. While creating real value for users, these companies are also capturing a disproportionate and expanding share of the value, and that's shaping our collective economic future. The very same technologies that promised to democratize business are now threatening to make it more monopolistic.

Beyond dominating individual markets, hub firms create and control essential connections in the networks that pervade our economy. Google's Android and related technologies form "competitive bottlenecks"; that is, they own access to billions of mobile consumers that other product and service providers want to reach. Google can not only exact a toll on transactions but also influence the flow of information and the data collected. Amazon's and Alibaba's marketplaces also connect vast numbers

of users with large numbers of retailers and manufacturers. Tencent's WeChat messaging platform aggregates a billion global users and provides a critical source of consumer access for businesses offering online banking, entertainment, transportation, and other services. The more users who join these networks, the more attractive (and even necessary) it becomes for enterprises to offer their products and services through them. By driving increasing returns to scale and controlling crucial competitive bottlenecks, these digital superpowers can become even mightier, extract disproportionate value, and tip the global competitive balance.

Hub firms don't compete in a traditional fashion—vying with existing products or services, perhaps with improved features or lower cost. Rather, they take the network-based assets that have already reached scale in one setting and then use them to enter another industry and "re-architect" its competitive structure—transforming it from product-driven to network-driven. They plug adjacent industries into the same competitive bottlenecks they already control.

For example, the Alibaba spin-off Ant Financial does not simply offer better payment services, a better credit card, or an improved investment management service; it builds on data from Alibaba's already vast user base to commoditize traditional financial services and reorganize a good chunk of the Chinese financial sector around the Ant Financial platform. The three-year-old service already has over half a billion users and plans to

expand well beyond China. Similarly, Google's automotive strategy does not simply entail creating an improved car; it leverages technologies and data advantages (many already at scale from billions of mobile consumers and millions of advertisers) to change the structure of the auto industry itself. (Disclosure: Both of us work or have worked with some of the firms mentioned in this article.)

If current trends continue, the hub economy will spread across more industries, further concentrating data, value, and power in the hands of a small number of firms employing a tiny fraction of the workforce. Disparity in firm valuation and individual wealth already causes widespread resentment. Over time, we can expect consumers, regulators, and even social movements to take an increasingly hostile stand against this concentration of value and economic connectivity. In a painfully ironic turn, after creating unprecedented opportunity across the global economy, digitization—and the trends it has given rise to—could exacerbate already dangerous levels of income inequality, undermine the economy, and even lead to social instability.

Idea in Brief

The Situation

A few digital superpowers, or hub firms, are capturing a disproportionate and growing share of the value being created in the global economy.

The Challenge

This trend threatens to exacerbate already dangerous levels of income inequality, undermine the economy, and destabilize society.

The Answer

While there are ways for companies that depend on hubs to defend their positions, the hubs themselves will have to do more to share economic value and sustain stakeholders.

Can these trends be reversed? We believe not. The "hub economy," as we will argue, is here to stay. But most companies will not become hubs, and they will need to respond astutely to the growing concentration of hub power. Digitizing operating capabilities will not be enough. Digital messaging platforms, for example, have already dealt a blow to telecom service providers; investment advisors still face threats from online financial-services companies. To remain competitive, companies will need to use their assets and capabilities differently, transform their core businesses, develop new revenue opportunities, and identify areas that can be defended from encroaching hub firms and others rushing in from previously disconnected economic sectors. Some companies have started on this path—Comcast, with its new Xfinity platform, is a notable example—but the majority, especially those in traditional sectors, still need to master the implications of network competition.

Most importantly, the very same hub firms that are transforming our economy must be part of the solution—and their leaders must step up. As Mark Zuckerberg articulated in his Harvard commencement address in May 2017, "we have a level

of wealth inequality that hurts everyone." Business as usual is not a good option. Witness the public concern about the roles that Facebook and Twitter played in the recent U.S. presidential election, Google's challenges with global regulatory bodies, criticism of Uber's culture and operating policies, and complaints that Airbnb's rental practices are racially discriminatory and harmful to municipal housing stocks, rents, and pricing.

Thoughtful hub strategies will create effective ways to share economic value, manage collective risks, and sustain the networks and communities we all ultimately depend on. If carmakers, major retailers, or media companies continue to go out of business, massive economic and social dislocation will ensue. And with governments and public opinion increasingly attuned to this problem, hub strategies that foster a more stable economy and united society will drive differentiation among the hub firms themselves.

We are encouraged by Facebook's response to the public outcry over "fake news"—hiring thousands of dedicated employees, shutting down tens of thousands of phony accounts, working with news sources to identify untrue claims, and offering guides for spotting false information. Similarly, Google's YouTube division invests in engineering, artificial intelligence, and human resources and collaborates with NGOs to ensure that videos promoting political extremists and terrorists are taken down promptly.

A real opportunity exists for hub firms to truly lead our economy. This will require hubs to fully consider the long-term societal impact of their decisions and to prioritize their ethical responsibilities to the large economic ecosystems that increasingly revolve around them. At the same time, the rest of us—whether in established enterprises or start-ups, in institutions or communities—will need to serve as checks and balances, helping to shape the hub economy by providing critical, informed input and, as needed, pushback.

The Digital Domino Effect

The emergence of economic hubs is rooted in three principles of digitization and network theory. The first is Moore's law, which states that computer processing power will double approximately every two years. The implication is that performance improvements will continue driving the augmentation and replacement of human activity with digital tools. This affects any industry that has integrated computers into its operations—which pretty much covers the entire economy. And advances in machine learning and cloud computing have only reinforced this trend.

The second principle involves connectivity. Most computing devices today have built-in network connectivity that allows them to communicate with one another. Modern digital technology enables the sharing of information at near-zero marginal cost, and digital networks are spreading rapidly.

Metcalfe's law states that a network's value increases with the number of nodes (connection points) or users—the dynamic we think of as network effects. This means that digital technology is enabling significant growth in value across our economy, particularly as open-network connections allow for the recombination of business offerings, such as the migration from payment tools to the broader financial services and insurance that we've seen at Ant Financial.

But while value is being created for everyone, value capture is getting more skewed and concentrated. This is because in networks, traffic begets more traffic, and as certain nodes become more heavily used, they attract additional attachments, which further increases their importance. This brings us to the third principle, a lesser-known dynamic originally posited by the physicist Albert-László Barabási: the notion that digital-network formation naturally leads to the emergence of positive feedback loops that create increasingly important, highly connected hubs. As digital networks carry more and more economic transactions, the economic power of network hubs, which connect consumers, firms, and even industries to one another, expands. Once a hub is highly connected (and enjoying increasing returns to scale) in one sector of the economy (such as mobile telecommunications), it will enjoy a crucial advantage as it begins to connect in a new sector (automobiles, for example). This can, in turn, drive more and more markets to tip, and the many players competing in traditionally separate industries get winnowed down to just a

few hub firms that capture a growing share of the overall economic value created—a kind of digital domino effect.

This phenomenon isn't new. But in recent years, the high degree of digital connectivity has dramatically sped up the transformation. Just a few years ago, cell phone manufacturers competed head-to-head for industry leadership in a traditional product market without appreciable network effects. Competition led to innovation and differentiation, with a business model delivering healthy profitability at scale for a dozen or so major competitors. But with the introduction of iOS and Android, the industry began to tip away from its hardware centricity to network structures centered on these multisided platforms. The platforms connected smartphones to a large number of apps and services. Each new app makes the platform it sits on more valuable, creating a powerful network effect that in turn creates a more daunting barrier to entry for new players. Today Motorola, Nokia, BlackBerry, and Palm are out of the mobile phone business, and Google and Apple are extracting the lion's share of the sector's value. The value captured by the large majority of complementors—the app developers and third-party manufacturers—is generally modest at best.

The domino effect is now spreading to other sectors and picking up speed. Music has already tipped to Apple, Google, and Spotify. E-commerce is following a similar path: Alibaba and Amazon are gaining more share and moving into traditional brick-and-mortar strongholds like groceries (witness Amazon's

acquisition of Whole Foods). We've already noted the growing power of WeChat in messaging and communications; along with Facebook and others, it's challenging traditional telecom service providers. On-premise computer and software offerings are losing ground to the cloud services provided by Amazon, Microsoft, Google, and Alibaba. In financial services, the big players are Ant, Paytm, Ingenico, and the independent start-up Wealthfront; in home entertainment, Amazon, Apple, Google, and Netflix dominate.

Where are powerful hub firms likely to emerge next? Health care, industrial products, and agriculture are three contenders. But let's examine how the digital domino effect could play out in another prime candidate, the automotive sector, which in the United States alone provides more than seven million jobs and generates close to a trillion dollars in yearly sales.

Re-architecting the Automotive Sector

As with many other products and services, cars are now connected to digital networks, essentially becoming rolling information and transaction nodes. This connectivity is reshaping the structure of the automotive industry. When cars were merely products, car sales were the main prize. But a new source of value is emerging: the connection to consumers in transit. Americans spend almost an hour, on average, getting to and from work every day, and commutes keep getting longer. Auto manufacturers, responding to consumer demand, have

already given hub firms access to dashboard screens in many cars; drivers can use Apple or Google apps on the car's built-in display instead of on their smartphones. If consumers embrace self-driving vehicles, that one hour of consumer access could be worth hundreds of billions of dollars in the U.S. alone.

Which companies will capitalize on the vast commercial potential of a new hour of free time for the world's car commuters? Hub firms like Alphabet and Apple are first in line. They already have bottleneck assets like maps and advertising networks at scale, and both are ready to create super-relevant ads pinpointed to the car's passengers and location. One logical add-on feature for autonomous vehicles would be a "Drive there" button that appears when an ad pops up (as already happens on Google's Waze app); pressing it would order the car to head to the touted destination.

In a future when people are no longer behind the wheel, cars will become less about the driving experience and more about the apps and services offered by automobiles as they ferry passengers around. Apart from a minority of cars actually driven for fun, differentiation will lessen, and the vehicle itself might well become commoditized. That will threaten manufacturers' core business: The car features that buyers will care most about—software and networks—will be largely outside the automakers' control, and their price premiums will go down.

The transformation will also upend a range of connected sectors—including insurance, automotive repairs and

maintenance, road construction, law enforcement, and infrastructure—as the digital dominos continue to fall. (See the exhibit "The connected-car ecosystem.")

The connected-car ecosystem

Three software platforms—Android Auto, Apple CarPlay, and, to a lesser extent, OpenCar—dominate the market for integrating smartphone functionality into vehicles. They constitute powerful bottleneck assets because they have scores of supply-chain partners (left) and they enable other stakeholders (right) to reach consumers. (Note: The companies, apps, and regulators listed are selected examples only.)

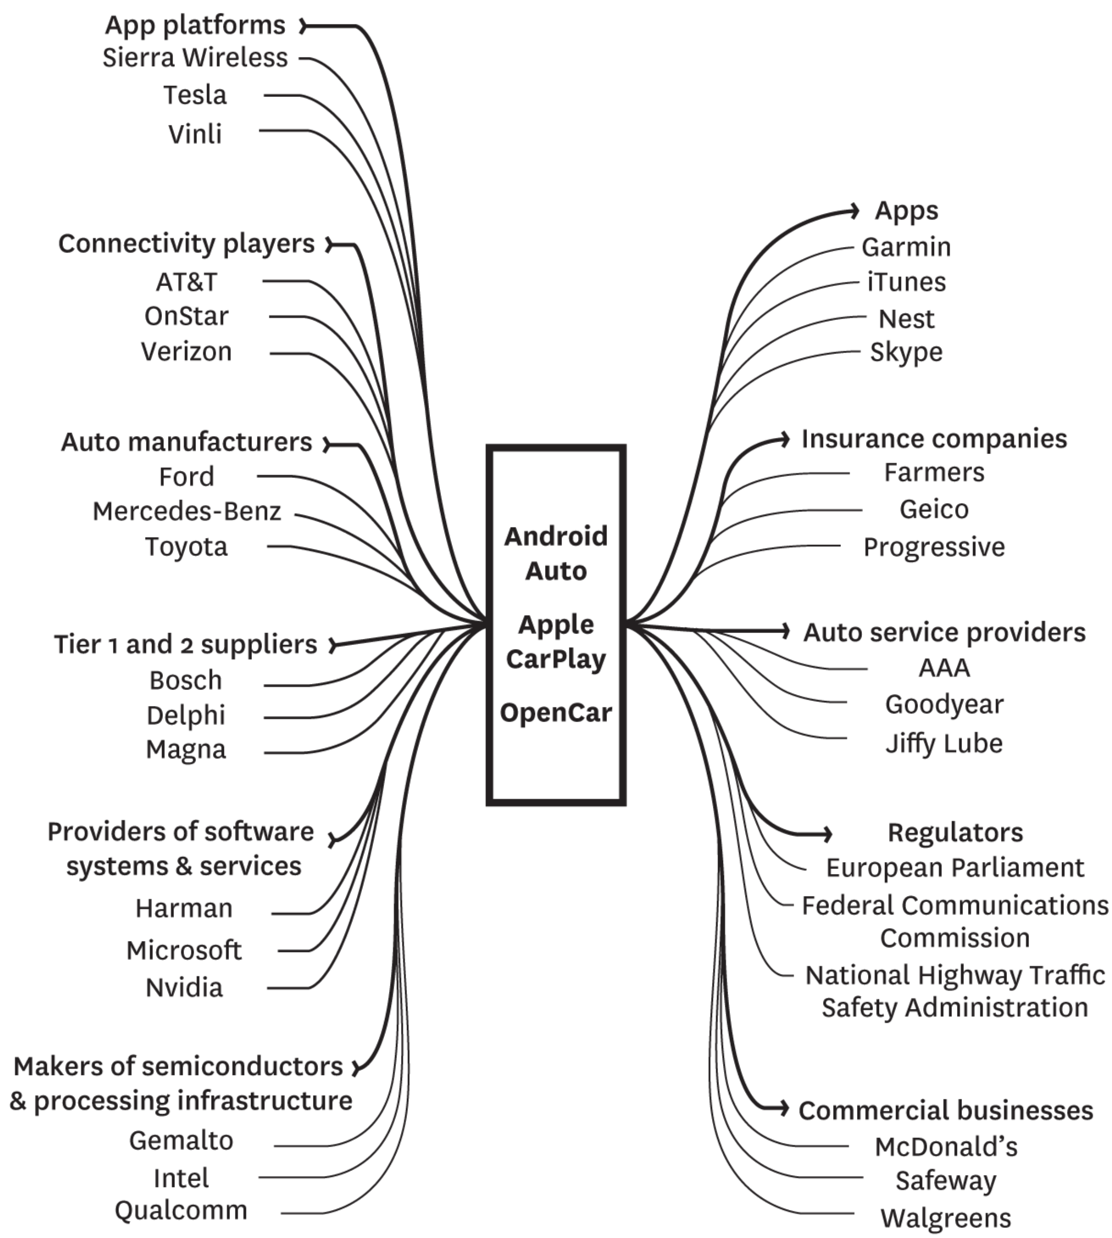

For existing auto manufacturers, the picture is grim but not hopeless. Some companies are exploring a pay-per-use model for their cars and are acquiring, launching, or partnering with car-as-a-service providers. GM, for one, invested $500 million in the ride-sharing service Lyft, and its luxury-car division is now

offering a monthly car subscription service. Daimler launched a car-sharing business called car2go. Several manufacturers have also invested in their own research into driverless vehicles or partnered with external providers.

Beyond these business-model experiments, automakers will need to play as the hubs do, by participating in the platform competition that will determine value capture in the sector. At least for the moment, alternatives to Google and Apple are scarce. One example is OpenCar, recently acquired by Inrix, a traditional auto supplier. Unlike Apple CarPlay and Google's Android Auto, which limit automaker-specific customization and require access to proprietary car data, the OpenCar framework is fully controlled by the car manufacturer. To take on the established giants, we believe that OpenCar and Inrix will have to develop an effective advertising or commerce platform or adopt some other indirect monetization strategy—and to do that, they'll probably need to partner with companies that have those capabilities.

To reach the scale required to be competitive, automotive companies that were once fierce rivals may need to join together. Here Technologies, which provides precision mapping data and location services, is an interesting example. Here has its roots in Navteq, one of the early online mapping companies, which was first bought by Nokia and later acquired by a consortium of Volkswagen, BMW, and Daimler (the multibillion-dollar price tag may have been too high for any single carmaker to stomach).

Here provides third-party developers with sophisticated tools and APIs for creating location-based ads and other services. The company represents an attempt by auto manufacturers to assemble a "federated" platform and, in doing so, neutralize the threat of a potential competitive bottleneck controlled by Google and Apple. The consortium could play a significant role in preventing automotive value capture from tipping completely toward existing hub firms.

Of course, successful collaboration depends on a common, strongly felt commitment. So as traditional enterprises position themselves for a fight, they must understand how the competitive dynamics in their industries have shifted.

Increasing Returns to Scale Are Hard to Beat

Competitive advantage in many industries is moderated by *decreasing* returns to scale. In traditional product and service businesses, the value creation curve typically flattens out as the number of consumers increases, as we see in the exhibit "Profiting from a growing customer base." A firm gains no particular advantage as its user base continues to increase beyond already efficient levels, which enables multiple competitors to coexist.

Profiting from a growing customer base

For traditional product and service businesses, gaining additional customers does not continue adding commensurate value after a certain point.

However, many platform businesses (Amazon, Facebook, and the like) become more and more valuable as more people and companies use them, connect with one another, and create network effects.

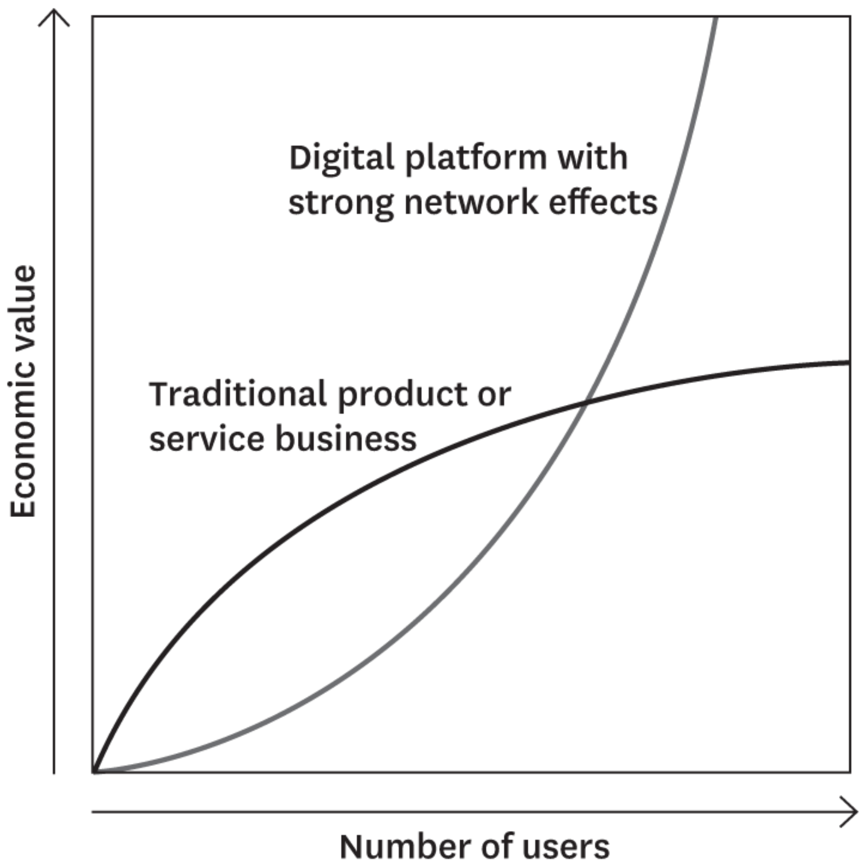

Some digital technologies, however, exhibit *increasing* returns to scale. A local advertising platform gets better and better as more and more users attract more and more ads. And as the number of ads increases, so does the ability to target the ads to the users, making individual ads more valuable. An advertising platform is thus similar to software platforms such as Windows, Linux, Android, and iOS, which exhibit increasing returns to scale—their growing value to consumers increases the number of

available apps, while the value to app developers rises as the number of consumers rises. The more consumers, the greater the incentive for developers to build apps, and the more apps there are, the more motivated consumers are to use their digital devices.

These considerations are important to the nature of hub competition. The economics of traditional decreasing returns make it possible for several competitors to coexist and provide differentiated value to attract users. That's the dynamic in the auto industry today, with many car manufacturers competing with one another to offer a variety of differentiated products. But the increasing returns in digital assets like ad platforms (or possibly driverless-car technology) will heighten the advantage of the competitor with the largest scale, the largest network of users, or the most data. And this is where the hub firms will most likely leverage their large and growing lead—and cause value to concentrate around them.

In contrast with traditional product and service businesses, network-based markets exhibiting increasing returns to scale will, over time, tip toward a narrow set of players. This implies that if a conventional decreasing-returns business (say, telecom or media) is threatened by a new type of competitor whose business model experiences increasing returns, the conventional player is in for a rough ride. With increasing returns to scale, a digital technology can provide a bottleneck to an entire

industrial sector. And left alone, competitive bottlenecks dramatically skew value capture away from traditional firms.

Pushing Back

Hub firms often compete against one another. Microsoft has made substantial investments in augmented reality in an effort to create a new hub and counterbalance the power that Google and Apple wield in the mobile space. Facebook acquired Oculus to force a similar structural shift in the emerging field of virtual reality. And a battle is looming in the smart-home arena, as Google, Apple, Microsoft, and Samsung attempt to reduce Amazon's early lead in voice-activated home technology.

But how does the rest of the economy deal with the increasing returns to scale of hub firms? With enough foresight and investment, traditional firms can resist by becoming hubs themselves, as we are seeing especially in the internet of things (IoT) space. GE is the classic example of this approach, with its investment in the Predix platform and the creation of GE Digital. [See the article "How I Remade GE," HBR, September–October 2017.] Other companies are following suit in different settings—for example, Verizon and Vodafone with their IoT platforms.

Firms can also shape competition by investing to ensure that there are multiple hubs in each sector—and even influencing which ones win. They can organize to support less-established platforms, thus making a particular hub more viable and an industry sector more competitive in the long term. Deutsche

Telekom, for instance, is partnering with Microsoft Azure (rather than Amazon Web Services) for cloud computing in Central Europe.

Most importantly, the value generated by networks will change as firms compete, innovate, and respond to community and regulatory pressure. Multihoming—a practice enabling participants on one hub's ecosystem to easily join another—can significantly mitigate the rise of hub power. For example, drivers and passengers routinely multihome across different ride-sharing platforms, often checking prices on Uber, Lyft, and Fasten to see which is offering the best deal. Retailers are starting to multihome across payment systems, supporting multiple solutions (such as Apple Pay, Google Wallet, and Samsung Pay). If multihoming is common, the market is less likely to tip to a single player, preserving competition and diffusing value capture. Indeed, companies will need to make their products and services available on multiple hubs and encourage the formation of new hubs to avoid being held hostage by one dominant player. Take the wireless-speaker manufacturer Sonos: It has ensured that its music system seamlessly integrates with as many music services as possible, including Apple Music, Amazon Music Unlimited, Google Play Music, Pandora, Spotify, and Tidal.

Collective action can also restructure economic networks, shape value creation and capture, and ease competitive bottlenecks. In the 1990s the open-source community organized to compete against Microsoft Windows with the Linux operating

system. That effort was actively supported by traditional players such as IBM and Hewlett-Packard and reinforced later by Google and Facebook. Today Linux (and Linux-related products) are firmly established in enterprises, consumer devices, and cloud computing. Similarly, the Mozilla open-source community and its Firefox browser broke Microsoft's grip on navigating the internet. Even Apple, notorious for its proprietary approach, relies on open-source software for its core operating systems and web services, and the infamous iPhone jailbreaking craze demonstrated both the extraordinary demand for third-party apps and the burgeoning supply of them.

Open source has grown beyond all expectations to create an increasingly essential legacy of common intellectual property, capabilities, and methodologies. Now collective action is going well beyond code sharing to include coordination on data aggregation, the use of common infrastructure, and the standardization of practices to further equilibrate the power of hubs. Efforts like OpenStreetMap are leading the way in maps, and Mozilla's Common Voice project is crowdsourcing global voice data to open up the speech-recognition bottleneck.

Collective action will be increasingly crucial to sustaining balance in the digital economy. As economic sectors coalesce into networks and as powerful hubs continue to form, other stakeholders will need to work together to ensure that hubs look after the interests of all network members. Cooperation will become more important for the rivals that orbit hubs; indeed,

strategic joint action by companies that are not hubs may be the best competitive antidote to the rising power of hub firms.

The public is also raising concerns about privacy, online tracking, cybersecurity, and data aggregation. Solutions being suggested include requirements for social network and data portability similar to the requirements for phone number portability that telecommunications regulators instituted to increase competition among phone service providers.

The Ethics of Network Leadership

The responsibility for sustaining our (digital) economy rests partly with the same leaders who are poised to control it. By developing such central positions of power and influence, hub firms have become de facto stewards of the long-term health of our economy. Leaders of hub companies need to realize that their organizations are analogous to “keystone” species in biological ecosystems—playing a critical role in maintaining their surroundings. Apple, Alibaba, Alphabet/Google, Amazon, and others that benefit disproportionately from the ecosystems they dominate have rational and ethical reasons to support the economic vitality of not just their direct participants but also the broader industries they serve. In particular, we argue that hub companies need to incorporate value *sharing* into their business models, along with value creation and value capture.

Building and maintaining a healthy ecosystem is in the best interests of hub companies. Amazon and Alibaba claim millions

of marketplace sellers, and they profit from every transaction those merchants make. Similarly, Google and Apple earn billions in revenue from the third-party apps that run on their platforms. Both companies already invest heavily in the developer community, providing programming frameworks, software tools, and opportunities and business models that enable developers to grow their businesses. But such efforts will need to be scaled up and refined as hub firms find themselves at the center of—and relying on—much larger and more-complex ecosystems. Preserving the strength and productivity of complementary communities should be a fundamental part of any hub firm's strategy.

Uber provides an interesting example of the repercussions of getting this wrong. Uber's viability depends on its relations with its drivers and riders, who have often criticized the company's practices. Under pressure from those communities—and from competitors that offer drivers the potential to earn more—Uber is making improvements. Still, its challenges suggest that no hub will maintain an advantage over the long term if it neglects the well-being of its ecosystem partners. Microsoft learned a hard lesson when it failed to maintain the health of its PC software ecosystem, losing out to the Linux community in cloud services.

But network ethics are not just about financial considerations; social concerns are equally important. Centralized platforms, such as Kiva for charitable impact investing and Airbnb for accommodation bookings, have been found to be susceptible to

racial discrimination. In Airbnb's case, external researchers convincingly demonstrated that African-American guests were especially likely to have their reservation requests rejected. The pressure is now on Airbnb to fight bias both by educating its proprietors and by modifying certain platform features. Additionally, as Airbnb continues to grow, it must work to ensure that its hosts heed municipal regulations, lest they face a potentially devastating regulatory backlash.

Indeed, if hubs do not promote the health and sustainability of the many firms and individuals in their networks, other forces will undoubtedly step in. Governments and regulators will increasingly act to encourage competition, protect consumer welfare, and foster economic stability. Consider the challenges Google faces in Europe, where regulators are concerned about the dominance of both its search advertising business and its Android platform.

The centralizing forces of digitization are not going to slow down anytime soon. The emergence of powerful hub firms is well under way, and the threats to global economic well-being are unmistakable. All actors in the economy—but particularly the hub firms themselves—should work to sustain the entire ecosystem and observe new principles, for both strategic and ethical reasons. Otherwise, we are all in serious trouble.

Originally published in September–October 2017. Reprint R1705F

The Leader's Guide to Corporate Culture

by Boris Groysberg, Jeremiah Lee, Jesse Price, and J. Yo-Jud Cheng

STRATEGY AND CULTURE ARE AMONG the primary levers at top leaders' disposal in their never-ending quest to maintain organizational viability and effectiveness. Strategy offers a formal logic for the company's goals and orients people around them. Culture expresses goals through values and beliefs and guides activity through shared assumptions and group norms.

Strategy provides clarity and focus for collective action and decision making. It relies on plans and sets of choices to mobilize people and can often be enforced by both concrete rewards for achieving goals and consequences for failing to do so. Ideally, it also incorporates adaptive elements that can scan and analyze the external environment and sense when changes are required to maintain continuity and growth. Leadership goes hand-in-hand with strategy formation, and most leaders understand the fundamentals. Culture, however, is a more elusive lever, because much of it is anchored in unspoken behaviors, mindsets, and social patterns.

For better *and* worse, culture and leadership are inextricably linked. Founders and influential leaders often set new cultures in motion and imprint values and assumptions that persist for decades. Over time an organization's leaders can also shape culture, through both conscious and unconscious actions (sometimes with unintended consequences). The best leaders we have observed are fully aware of the multiple cultures within which they are embedded, can sense when change is required, and can deftly influence the process.

Unfortunately, in our experience it is far more common for leaders seeking to build high-performing organizations to be confounded by culture. Indeed, many either let it go unmanaged or relegate it to the HR function, where it becomes a secondary concern for the business. They may lay out detailed, thoughtful plans for strategy and execution, but because they don't understand culture's power and dynamics, their plans go off the rails. As someone once said, culture eats strategy for breakfast.

It doesn't have to be that way. Our work suggests that culture can, in fact, be managed. The first and most important step leaders can take to maximize its value and minimize its risks is to become fully aware of how it works. By integrating findings from more than 100 of the most commonly used social and behavioral models, we have identified eight styles that distinguish a culture and can be measured. (We gratefully acknowledge the rich history of cultural studies—going all the way back to the earliest explorations of human nature—on which our work builds.) Using

this framework, leaders can model the impact of culture on their business and assess its alignment with strategy. We also suggest how culture can help them achieve change and build organizations that thrive in even the most trying times.

Defining Culture

Culture is the tacit social order of an organization: It shapes attitudes and behaviors in wide-ranging and durable ways. Cultural norms define what is encouraged, discouraged, accepted, or rejected within a group. When properly aligned with personal values, drives, and needs, culture can unleash tremendous amounts of energy toward a shared purpose and foster an organization's capacity to thrive.

Culture can also evolve flexibly and autonomously in response to changing opportunities and demands. Whereas strategy is typically determined by the C-suite, culture can fluidly blend the intentions of top leaders with the knowledge and experiences of frontline employees.

Idea in Brief

Executives are often confounded by culture, because much of it is anchored in unspoken behaviors, mindsets, and social patterns. Many leaders either let it go unmanaged or relegate it to HR, where it becomes a secondary concern for the business. This is a mistake, because properly managed, culture can help them achieve change and build organizations that will thrive in even the most trying times.

The authors have reviewed the literature on culture and distilled eight distinct culture styles: *caring*, focused on relationships and mutual trust; *purpose*, exemplified by idealism and altruism; *learning*, characterized by exploration, expansiveness, and creativity; *enjoyment*, expressed through fun and excitement; *results*, characterized by achievement and winning; *authority*, defined by strength, decisiveness, and boldness; *safety*, defined by planning, caution, and preparedness; and *order*, focused on respect, structure, and shared norms.

These eight styles fit into an "integrated culture framework" according to the degree to which they reflect independence or interdependence (people interactions) and flexibility or stability (response to change). They can be used to diagnose and describe highly complex and diverse behavioral patterns in a culture and to model how likely an individual leader is to align with and shape that culture.

Through research and practical experience, the authors have arrived at five insights regarding culture's effect on companies' success: (1) When aligned with strategy and leadership, a strong culture drives positive organizational outcomes. (2) Selecting or developing leaders for the future requires a forward-looking strategy and culture. (3) In a merger, designing a new culture on the basis of complementary strengths can speed up integration and create more value over time. (4) In a dynamic, uncertain environment, in which organizations must be more agile, learning gains importance. (5) A strong culture can be a significant liability when it is misaligned with strategy.

The academic literature on the subject is vast. Our review of it revealed many formal definitions of organizational culture and a variety of models and methods for assessing it. Numerous processes exist for creating and changing it. Agreement on specifics is sparse across these definitions, models, and methods, but through a synthesis of seminal work by Edgar Schein,

Shalom Schwartz, Geert Hofstede, and other leading scholars, we have identified four generally accepted attributes:

Shared

Culture is a group phenomenon. It cannot exist solely within a single person, nor is it simply the average of individual characteristics. It resides in shared behaviors, values, and assumptions and is most commonly experienced through the norms and expectations of a group—that is, the unwritten rules.

Pervasive

Culture permeates multiple levels and applies very broadly in an organization; sometimes it is even conflated with the organization itself. It is manifest in collective behaviors, physical environments, group rituals, visible symbols, stories, and legends. Other aspects of culture are unseen, such as mindsets, motivations, unspoken assumptions, and what David Rooke and William Torbert refer to as "action logics" (mental models of how to interpret and respond to the world around you).

Enduring

Culture can direct the thoughts and actions of group members over the long term. It develops through critical events in the collective life and learning of a group. Its endurance is explained in part by the attraction-selection-attrition model first introduced by Benjamin Schneider: People are drawn to organizations with characteristics similar to their own;

organizations are more likely to select individuals who seem to "fit in"; and over time those who don't fit in tend to leave. Thus culture becomes a self-reinforcing social pattern that grows increasingly resistant to change and outside influences.

Implicit

An important and often overlooked aspect of culture is that despite its subliminal nature, people are effectively hardwired to recognize and respond to it instinctively. It acts as a kind of silent language. Shalom Schwartz and E.O. Wilson have shown through their research how evolutionary processes shaped human capacity; because the ability to sense and respond to culture is universal, certain themes should be expected to recur across the many models, definitions, and studies in the field. That is exactly what we have discovered in our research over the past few decades.

Eight Distinct Culture Styles

Our review of the literature for commonalities and central concepts revealed two primary dimensions that apply regardless of organization type, size, industry, or geography: people interactions and response to change. Understanding a company's culture requires determining where it falls along these two dimensions.

People interactions

An organization's orientation toward people interactions and coordination will fall on a spectrum from highly independent to highly interdependent. Cultures that lean toward the former place greater value on autonomy, individual action, and competition. Those that lean toward the latter emphasize integration, managing relationships, and coordinating group effort. People in such cultures tend to collaborate and to see success through the lens of the group.

Response to change

Whereas some cultures emphasize stability—prioritizing consistency, predictability, and maintenance of the status quo—others emphasize flexibility, adaptability, and receptiveness to change. Those that favor stability tend to follow rules, use control structures such as seniority-based staffing, reinforce hierarchy, and strive for efficiency. Those that favor flexibility tend to prioritize innovation, openness, diversity, and a longer-term orientation. (Kim Cameron, Robert Quinn, and Robert Ernest are among the researchers who employ similar dimensions in their culture frameworks.)

By applying this fundamental insight about the dimensions of people interactions and response to change, we have identified eight styles that apply to both organizational cultures and individual leaders. Researchers at Spencer Stuart (including two of this article's authors) have interdependently studied and refined this list of styles across both levels over the past two decades.

Caring focuses on relationships and mutual trust. Work environments are warm, collaborative, and welcoming places where people help and support one another. Employees are united by loyalty; leaders emphasize sincerity, teamwork, and positive relationships.

Purpose is exemplified by idealism and altruism. Work environments are tolerant, compassionate places where people try to do good for the long-term future of the world. Employees are united by a focus on sustainability and global communities; leaders emphasize shared ideals and contributing to a greater cause.

Learning is characterized by exploration, expansiveness, and creativity. Work environments are inventive and open-minded places where people spark new ideas and explore alternatives. Employees are united by curiosity; leaders emphasize innovation, knowledge, and adventure.

Enjoyment is expressed through fun and excitement. Work environments are lighthearted places where people tend to do what makes them happy. Employees are united by playfulness and stimulation; leaders emphasize spontaneity and a sense of humor.

Results is characterized by achievement and winning. Work environments are outcome-oriented and merit-based places where people aspire to achieve top performance. Employees are united by a drive for capability and success; leaders emphasize goal accomplishment.

Authority is defined by strength, decisiveness, and boldness. Work environments are competitive places where people strive to gain personal advantage. Employees are united by strong control; leaders emphasize confidence and dominance.

Safety is defined by planning, caution, and preparedness. Work environments are predictable places where people are risk-conscious and think things through carefully. Employees are united by a desire to feel protected and anticipate change; leaders emphasize being realistic and planning ahead.

Order is focused on respect, structure, and shared norms. Work environments are methodical places where people tend to play by the rules and want to fit in. Employees are united by cooperation; leaders emphasize shared procedures and time-honored customs.

These eight styles fit into our integrated culture framework (see the exhibit "Integrated culture: The framework") according to the degree to which they reflect independence or interdependence (people interactions) and flexibility or stability (response to change). Styles that are adjacent in the framework, such as *safety* and *order*, frequently coexist within organizations and their people. In contrast, styles that are located across from each other, such as *safety* and *learning*, are less likely to be found together and require more organizational energy to maintain simultaneously. Each style has advantages and disadvantages, and no style is inherently better than another. An organizational culture can be defined by the absolute and relative strengths of

each of the eight and by the degree of employee agreement about which styles characterize the organization. A powerful feature of this framework, which differentiates it from other models, is that it can also be used to define individuals' styles and the values of leaders and employees.

Integrated culture: The framework

On the basis of decades of experience analyzing organizations, executives, and employees, we developed a rigorous, comprehensive model to identify the key attributes of both group culture and individual leadership styles. Eight characteristics emerge when we map cultures along two dimensions: how people interact (independence to interdependence) and their response to change (flexibility to stability). The relative salience of these eight styles differs across organizations, though nearly all are strongly characterized by results *and* caring.

The spatial relationships are important. Proximate styles, such as safety *and* order, *or* learning *and* enjoyment, *will coexist more easily than styles that are far apart on the chart, such as* authority *and* purpose, *or* safety *and* learning. *Achieving a culture of* authority *often means gaining the advantages (and living with the disadvantages) of that culture but missing out on the advantages (and avoiding the disadvantages) of a culture of* purpose.

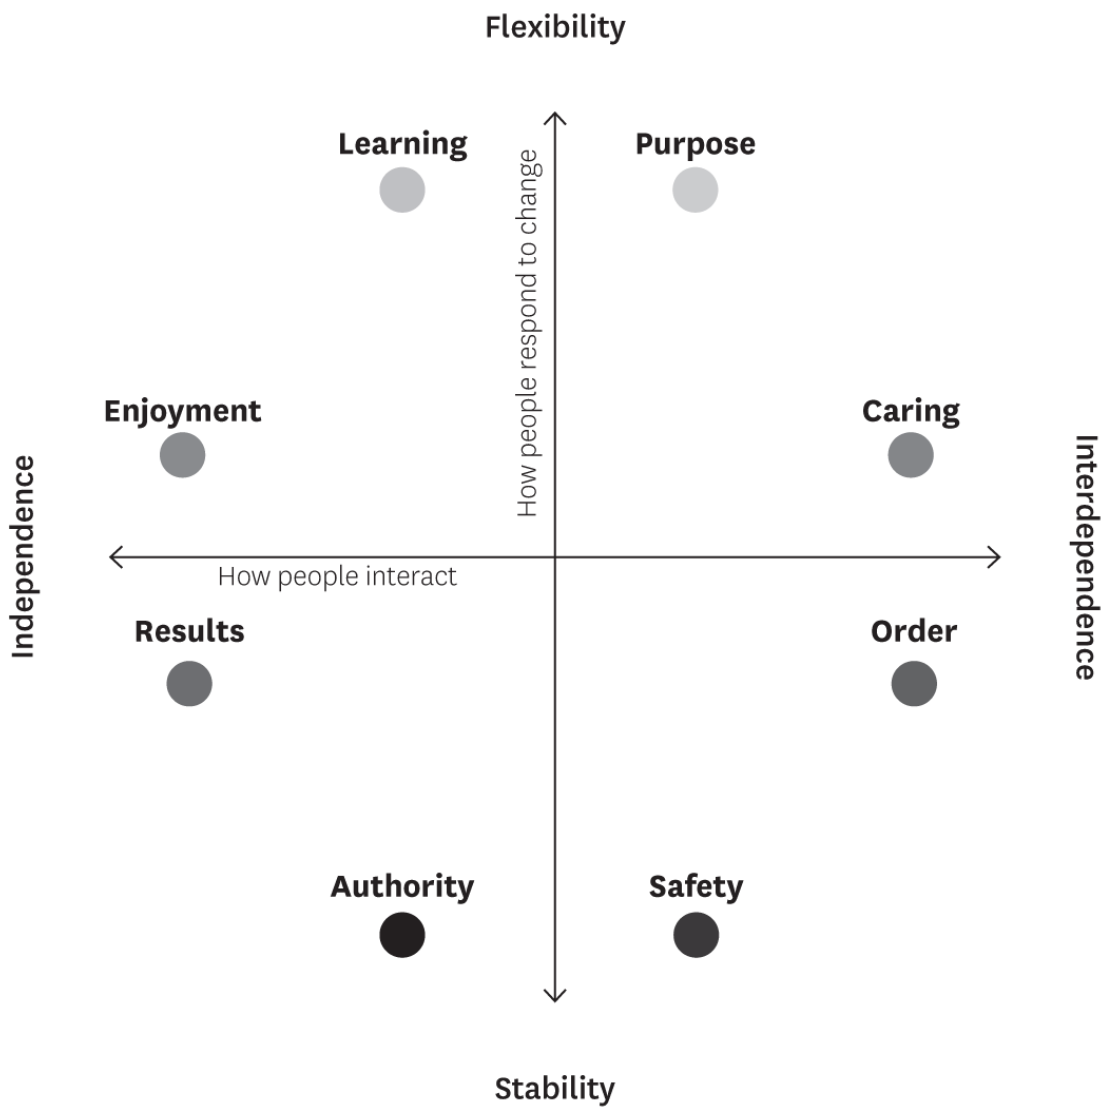

Source: Spencer Stuart.

Inherent in the framework are fundamental trade-offs. Although each style can be beneficial, natural constraints and competing demands force difficult choices about which values to emphasize and how people are expected to behave. It is common to find organizations with cultures that emphasize both *results* and *caring*, but this combination can be confusing to employees. Are they expected to optimize individual goals and strive for

outcomes at all costs, or should they work as a team and emphasize collaboration and shared success? The nature of the work itself, the business strategy, or the design of the organization may make it difficult for employees to be equally *results* focused and *caring*.

Integrated culture: Leader statements

Top leaders and founders often express cultural sentiments within the public domain, either intentionally or unintentionally. Such statements can provide important clues to how these leaders are thinking about and leading their organizations' cultures.

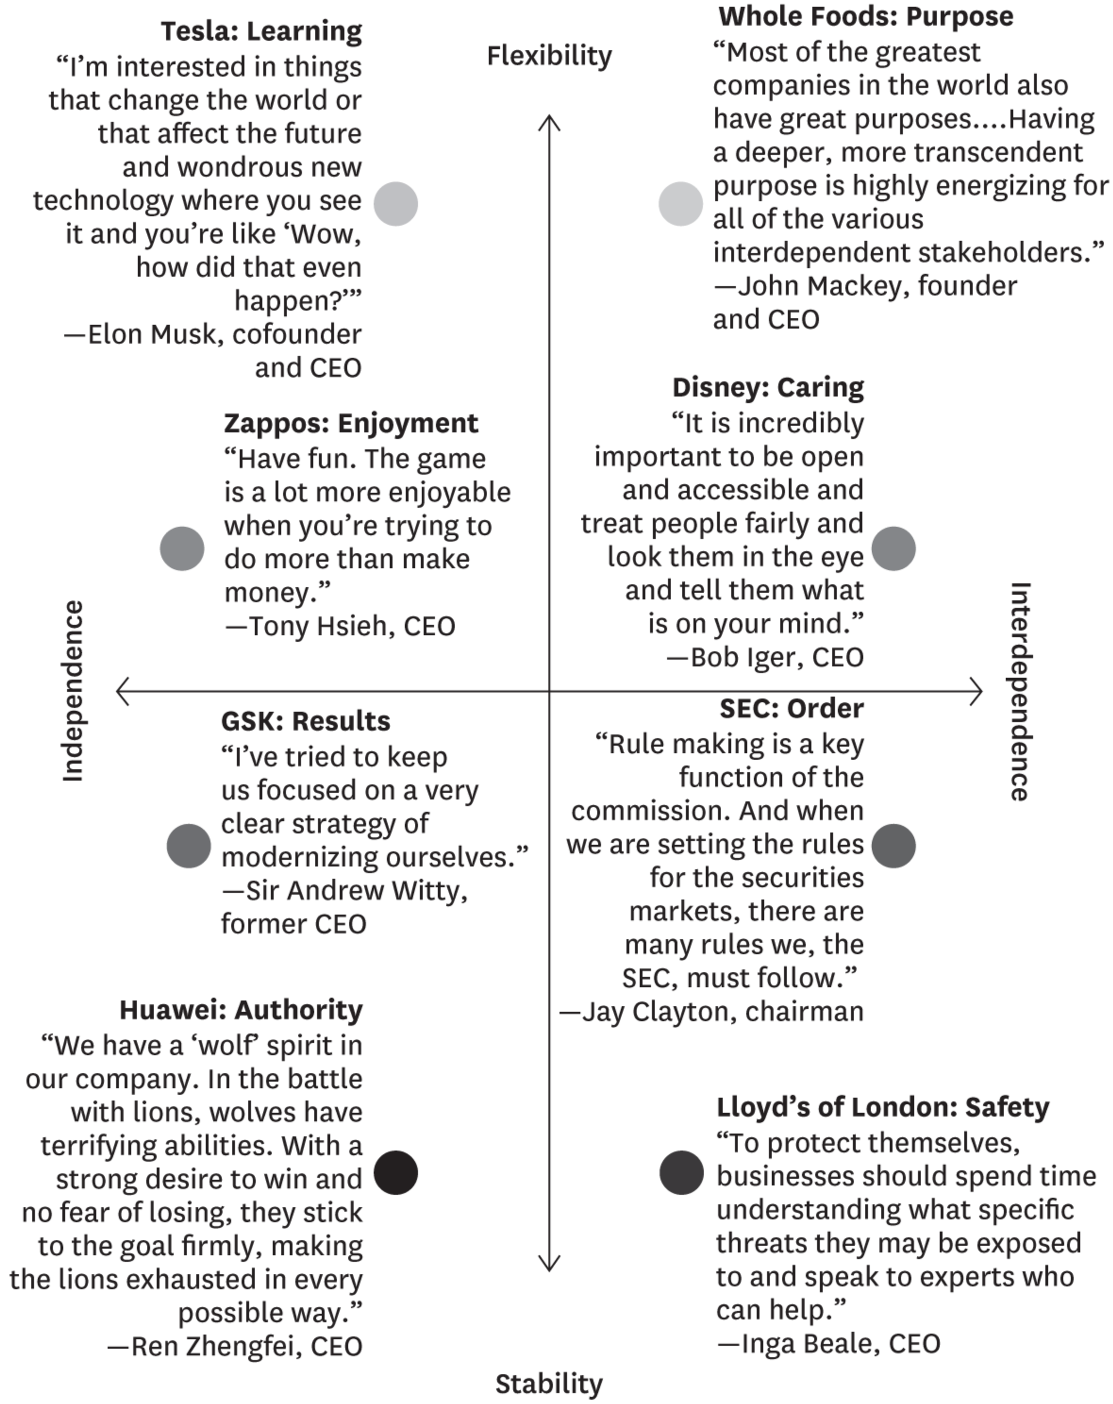

The pros and cons of culture styles

Every culture style has strengths and weaknesses. The table below summarizes the advantages and disadvantages of each style and how frequently it appears as a defining culture characteristic among the companies in our study.

Culture style	Advantages	Disadvantages	Ranked 1st or 2nd
Caring Warm, sincere, relational	Improved teamwork, engagement, communication, trust, and sense of belonging	Overemphasis on consensus building may reduce exploration of options, stifle competitiveness, and slow decision making	63%
Purpose Purpose driven, idealistic, tolerant	Improved appreciation for diversity, sustainability, and social responsibility	Overemphasis on a long-term purpose and ideals may get in the way of practical and immediate concerns	9%
Learning Open, inventive, exploring	Improved innovation, agility, and organizational learning	Overemphasis on exploration may lead to a lack of focus and an inability to exploit existing advantages	7%
Enjoyment Playful, instinctive, fun loving	Improved employee morale, engagement, and creativity	Overemphasis on autonomy and engagement may lead to a lack of discipline and create possible compliance or governance issues	2%
Results Achievement driven, goal focused	Improved execution, external focus, capability building, and goal achievement	Overemphasis on achieving results may lead to communication and collaboration breakdowns and higher levels of stress and anxiety	89%
Authority Bold, decisive, dominant	Improved speed of decision making and responsiveness to threats or crises	Overemphasis on strong authority and bold decision making may lead to politics, conflict, and a psychologically unsafe work environment	4%
Safety Realistic, careful, prepared	Improved risk management, stability, and business continuity	Overemphasis on standardization and formalization may lead to bureaucracy, inflexibility, and dehumanization of the work environment	8%
Order Rule abiding, respectful, cooperative	Improved operational efficiency, reduced conflict, and greater civic-mindedness	Overemphasis on rules and traditions may reduce individualism, stifle creativity, and limit organizational agility	15%

Note: Sum of percentages is greater than 100 because styles were counted as dominant if they were ranked 1 or 2 overall.

In contrast, a culture that emphasizes *caring* and *order* encourages a work environment in which teamwork, trust, and respect are paramount. The two styles are mutually reinforcing, which can be beneficial but can also present challenges. The benefits are strong loyalty, retention of talent, lack of conflict, and high levels of engagement. The challenges are a tendency toward groupthink, reliance on consensus-based decisions,

avoidance of difficult issues, and a calcified sense of "us versus them." Leaders who are more focused on *results* and *learning* may find the combination of *caring* and *order* stifling when they seek to drive entrepreneurship and change. Savvy leaders make use of existing cultural strengths and have a nuanced understanding of how to initiate change. They might rely on the participative nature of a culture focused on *caring* and *order* to engage team members and simultaneously identify a *learning*-oriented "insider" who has the trust of his or her peers to advocate for change through relationship networks.

The eight styles can be used to diagnose and describe highly complex and diverse behavioral patterns in a culture and to model how likely an individual leader is to align with and shape that culture. Using this framework and multilevel approach, managers can:

- Understand their organization's culture and assess its intended and unintended effects
- Evaluate the level of consistency in employees' views of the culture
- Identify subcultures that may account for higher or lower group performance
- Pinpoint differences between legacy cultures during mergers and acquisitions
- Rapidly orient new executives to the culture they are joining and help them determine the most effective way to

lead employees

- Measure the degree of alignment between individual leadership styles and organizational culture to determine what impact a leader might have
- Design an aspirational culture and communicate the changes necessary to achieve it

The Link Between Culture and Outcomes

Our research and practical experience have shown that when you are evaluating how culture affects outcomes, the context in which the organization operates—geographic region, industry, strategy, leadership, and company structure—matters, as does the strength of the culture. (See "Context, Conditions, and Culture," the sidebar at the end of this article.) What worked in the past may no longer work in the future, and what worked for one company may not work for another.

We have arrived at the following insights:

When aligned with strategy and leadership, a strong culture drives positive organizational outcomes

Consider the case of a best-in-class retailer headquartered in the United States. The company had viewed its first priority as providing top-notch customer service. It accomplished this with a simple rule—Do right by the customer—that encouraged employees to use their judgment when providing service. A core HR training practice was to help every salesperson see customer

interactions as an opportunity to create “service stories that become legendary.” Employees were reminded to define service from the customer’s perspective, to constantly engage customers with questions geared toward understanding their specific needs and preferences, and to go beyond their expectations.

In measuring the culture of this company, we found that like many other large retailers, it was characterized primarily by a combination of *results* and *caring*. Unlike many other retailers, however, it had a culture that was also very flexible, *learning* oriented, and focused on *purpose*. As one top executive explained, “We have freedom as long as we take good care of the customer.”

Furthermore, the company’s values and norms were very clear to everyone and consistently shared throughout the organization. As the retailer expanded into new segments and geographies over the years, the leadership strove to maintain an intense customer focus without diluting its cherished culture. Although the company had historically focused on developing leaders from within—who were natural culture carriers—recruiting outsiders became necessary as it grew. The company preserved its culture through this change by carefully assessing new leaders and designing an onboarding process that reinforced core values and norms.

Culture is a powerful differentiator for this company because it is strongly aligned with strategy and leadership. Delivering outstanding customer service requires a culture and a mindset

that emphasize achievement, impeccable service, and problem solving through autonomy and inventiveness. Not surprisingly, those qualities have led to a variety of positive outcomes for the company, including robust growth and international expansion, numerous customer service awards, and frequent appearances on lists of the best companies to work for.

Selecting or developing leaders for the future requires a forward-looking strategy and culture

The chief executive of an agriculture business was planning to retire, spurring rumors about a hostile takeover. The CEO was actively grooming a successor, an insider who had been with the company for more than 20 years. Our analysis revealed an organizational culture that strongly emphasized *caring* and *purpose*. As one leader reflected, “You feel like part of a large family when you become an employee at this company.”

The potential successor understood the culture but was far more risk-averse (*safety*) and respectful of traditions (*order*) than the rest of the company. Given the takeover rumors, top leaders and managers told the CEO that they believed the company needed to take a more aggressive and action-oriented stance in the future. The board decided to consider the internal candidate alongside people from outside the company.

Three external candidates emerged: one who was aligned with the current culture (*purpose*), one who would be a risk taker and innovative (*learning*), and one who was hard-driving and competitive (*authority*). After considerable deliberation, the

board chose the highly competitive leader with the *authority* style. Soon afterward an activist investor attempted a hostile takeover, and the new CEO was able to navigate through the precarious situation, keep the company independent, and simultaneously begin to restructure in preparation for growth.

In a merger, designing a new culture on the basis of complementary strengths can speed up integration and create more value over time

Mergers and acquisitions can either create or destroy value. Numerous studies have shown that cultural dynamics represent one of the greatest yet most frequently overlooked determinants of integration success and postmerger performance.

For example, senior leaders from two merging international food retailers had invested heavily in their organizations' cultures and wanted to preserve their unique strengths and distinct heritages. An assessment of the cultures revealed shared values and areas of compatibility that could provide a foundation for the combined culture, along with important differences for which leaders would have to plan: Both companies emphasized *results*, *caring*, and *order* and valued high-quality food, good service, treating employees fairly, and maintaining a local mindset. But one operated in a more top-down manner and scored much higher on *authority*, especially in the behavior of leaders.

Because both companies valued teamwork and investments in the local community, the leaders prioritized *caring* and *purpose*.

At the same time, their strategy required that they shift from top-down *authority* to a *learning* style that would encourage innovation in new-store formats and online retailing. As one senior leader said of the strategic aspiration, "We need to dare to do things differently, not play by the old rule books."

Once they had agreed on a culture, a rigorous assessment process identified leaders at both organizations whose personal style and values would allow them to serve as bridges to and champions for it. Then a program was launched to promote cultural alignment within 30 top teams, with an emphasis on clarifying priorities, making authentic connections, and developing team norms that would bring the new culture to life.

Finally, structural elements of the new organization were redesigned with culture in mind. A model for leadership was developed that encompassed recruitment, talent assessment, training and development, performance management, reward systems, and promotions. Such design considerations are often overlooked during organizational change, but if systems and structures don't align with cultural and leadership imperatives, progress can be derailed.

In a dynamic, uncertain environment, in which organizations must be more agile, *learning* gains importance

It's not surprising that *results* is the most common culture style among all the companies we have studied. Yet during a decade of helping leaders design aspirational cultures, we have seen a clear trend toward prioritizing *learning* to promote innovation

and agility as businesses respond to increasingly less predictable and more complex environments. And although *learning* ranks fourth within our broader database, small companies (200 employees or fewer) and those in newer industries (such as software, technology, and wireless equipment) accord it higher values.

Consider one Silicon Valley–based technology company we worked with. Though it had built a strong business and invested in unique technology and top engineering talent, its revenue growth was starting to decline as newer, nimbler competitors made strides in a field exploding with innovation and business model disruption. Company leaders viewed the culture as a differentiator for the business and decided to diagnose, strengthen, and evolve it. We found a culture that was intensely *results* focused, team based (*caring*), and exploratory (a combination of *enjoyment* and *learning*).

After examining the overall business strategy and gaining input from employees, leaders aimed for a culture that was even more focused on *learning* and adopted our framework as a new language for the organization in its daily work. They initiated conversations between managers and employees about how to emphasize innovation and exploration. Although it takes time to change a culture, we found that the company had made notable progress just one year later. And even as it prepared for an impending sale amid ever greater competition and consolidation, employee engagement scores were on the rise.

A strong culture can be a significant liability when it is misaligned with strategy

We studied a Europe-based industrial services organization whose industry began to experience rapid and unprecedented changes in customer expectations, regulatory demands, and competitive dynamics. The company's strategy, which had historically emphasized cost leadership, needed to shift toward greater service differentiation in response. But its strong culture presented a roadblock to success.

We diagnosed the culture as highly *results* oriented, *caring*, and *order* seeking, with a top-down emphasis on *authority*. The company's leaders decided to shape it to be much more *purpose*-driven, enabling, open, and team based, which would entail an increase in *caring* along with *learning* and *purpose* and a decrease in *authority* and *results*.

This shift was particularly challenging because the current culture had served the organization well for many years, while the industry emphasized efficiency and *results*. Most managers still viewed it as a strength and fought to preserve it, threatening success for the new strategic direction.

Cultural change is daunting for any organization, but as this company realized, it's not impossible. The CEO introduced new leadership development and team coaching programs and training opportunities that would help leaders feel more comfortable with cultural evolution. When people departed, the company carefully selected new leaders who would provide supporting values, such as *caring*, and increased the emphasis on

a shared *purpose*. The benefits of this strategic and cultural shift took the form of an increasingly diverse array of integrated service offerings and strong growth, particularly in emerging markets.

Four Levers for Evolving a Culture

Unlike developing and executing a business plan, changing a company's culture is inextricable from the emotional and social dynamics of people in the organization. We have found that four practices in particular lead to successful culture change:

Articulate the aspiration

Much like defining a new strategy, creating a new culture should begin with an analysis of the current one, using a framework that can be openly discussed throughout the organization. Leaders must understand what outcomes the culture produces and how it does or doesn't align with current and anticipated market and business conditions. For example, if the company's primary culture styles are *results* and *authority* but it exists in a rapidly changing industry, shifting toward *learning* or *enjoyment* (while maintaining a focus on *results*) may be appropriate.

An aspirational culture suggests the high-level principles that guide organizational initiatives, as at the technology company that sought to boost agility and flexibility amid increasing competition. Change might be framed in terms of real and present business challenges and opportunities as well as

aspirations and trends. Because of culture's somewhat ambiguous and hidden nature, referring to tangible problems, such as market pressures or the challenges of growth, helps people better understand and connect to the need for change.

Select and develop leaders who align with the target culture

Leaders serve as important catalysts for change by encouraging it at all levels and creating a safe climate and what Edgar Schein calls "practice fields." Candidates for recruitment should be evaluated on their alignment with the target. A single model that can assess both organizational culture and individual leadership styles is critical for this activity.

Incumbent leaders who are unsupportive of desired change can be engaged and re-energized through training and education about the important relationship between culture and strategic direction. Often they will support the change after they understand its relevance, its anticipated benefits, and the impact that they personally can have on moving the organization toward the aspiration. However, culture change can and does lead to turnover: Some people move on because they feel they are no longer a good fit for the organization, and others are asked to leave if they jeopardize needed evolution.

Use organizational conversations about culture to underscore the importance of change

To shift the shared norms, beliefs, and implicit understandings within an organization, colleagues can talk one another through

the change. Our integrated culture framework can be used to discuss current and desired culture styles and also differences in how senior leaders operate. As employees start to recognize that their leaders are talking about new business outcomes—innovation instead of quarterly earnings, for example—they will begin to behave differently themselves, creating a positive feedback loop.

Various kinds of organizational conversations, such as road shows, listening tours, and structured group discussion, can support change. Social media platforms encourage conversations between senior managers and frontline employees. Influential change champions can advocate for a culture shift through their language and actions. The technology company made a meaningful change in its culture and employee engagement by creating a structured framework for dialogue and cultivating widespread discussion.

Reinforce the desired change through organizational design

When a company's structures, systems, and processes are aligned and support the aspirational culture and strategy, instigating new culture styles and behaviors will become far easier. For example, performance management can be used to encourage employees to embody aspirational cultural attributes. Training practices can reinforce the target culture as the organization grows and adds new people. The degree of centralization and the number of hierarchical levels in the organizational structure can be adjusted to reinforce behaviors inherent to the aspirational

culture. Leading scholars such as Henry Mintzberg have shown how organizational structure and other design features can have a profound impact over time on how people think and behave within an organization.

Putting It All Together

All four levers came together at a traditional manufacturer that was trying to become a full solutions provider. The change started with reformulating the strategy and was reinforced by a major brand campaign. But the president understood that the company's culture represented the biggest barrier to change and that the top leaders were the greatest lever for evolving the culture.

The culture was characterized by a drive for *results* followed by *caring* and *purpose*, the last of which was unusually strong for the industry. One employee described the company as "a talented and committed group of people focused on doing good for the planet, with genuine desire, support, and encouragement to make a difference in the community." Whereas the broader culture was highly collaborative, with flat decision making, leaders were seen as top-down, hierarchical, and sometimes political, which discouraged risk taking.

The top leaders reviewed their culture's strengths and the gaps in their own styles and discussed what was needed to achieve their strategic aspirations. They agreed that they needed more risk taking and autonomy and less hierarchy and centralized

decision making. The president restructured the leadership team around strong business line leaders, freeing up time to become a better advocate for the culture and to focus more on customers.

The top team then invited a group of 100 middle managers into the conversation through a series of biannual leadership conferences. The first one established a platform for input, feedback, and the cocreation of an organizational change plan with clear cultural priorities. The president organized these managers into teams focused on critical business challenges. Each team was required to go outside the company to source ideas, to develop solutions, and to present its findings to the group for feedback. This initiative placed middle managers in change roles that would traditionally have been filled by vice presidents, giving them greater autonomy in fostering a *learning*-based culture. The intent was to create real benefits for the business while evolving the culture.

The president also initiated a program to identify employees who had positive disruptive ideas and working styles. These people were put on project teams that addressed key innovation priorities. The teams immediately began improving business results, both in core commercial metrics and in culture and engagement. After only one year employee engagement scores jumped a full 10 points, and customer Net Promoter Scores reached an all-time high—providing strong client references for the company's new and innovative solutions.

It is possible—in fact, vital—to improve organizational performance through culture change, using the simple but powerful models and methods in this article. First leaders must become aware of the culture that operates in their organization. Next they can define an aspirational target culture. Finally they can master the core change practices of articulation of the aspiration, leadership alignment, organizational conversation, and organizational design. Leading with culture may be among the few sources of sustainable competitive advantage left to companies today. Successful leaders will stop regarding culture with frustration and instead use it as a fundamental management tool.

What's Your Organization's Cultural Profile?

Before you begin an initiative to shape your organization's culture, it's important to explore where it is today. This worksheet and the questions that follow can help you formulate a preliminary assessment of your culture and get the conversation started.

Consider how your organization currently operates, what is valued, how people behave, and what unifies them. Partner with a colleague and independently rate each statement according to how well it describes your organization.

Add the two ratings in each row and then rank the eight styles. The higher the total, the stronger the match.

Compare your rankings with your colleague's and discuss the following questions:

- What do you like most about the current culture?
- What behaviors and mindsets might you evolve?
- How effective are your organization's leaders at role modeling the culture?
- What are the characteristics of people who are most successful in your culture?
- When new people don't succeed in your culture, what is the most common reason?

On a scale of 1-5, rate how well each of these statements describes your organization

1 = Not at all well 2 = Not very well 3 = Somewhat well 4 = Very well 5 = Extremely well

The organization is focused on:					The organization feels like:					Total
Collaboration and mutual trust					A big family					Caring
1	2	3	4	5	1	2	3	4	5	
Compassion and tolerance					An idealistic community or cause					Purpose
1	2	3	4	5	1	2	3	4	5	
Exploration and creativity					A dynamic project					Learning
1	2	3	4	5	1	2	3	4	5	
Fun and excitement					A celebration					Enjoyment
1	2	3	4	5	1	2	3	4	5	
Achievement and winning					A meritocracy					Results
1	2	3	4	5	1	2	3	4	5	
Strength and boldness					A competitive arena					Authority
1	2	3	4	5	1	2	3	4	5	
Planning and caution					A meticulously planned operation					Safety
1	2	3	4	5	1	2	3	4	5	
Structure and stability					A smoothly running machine					Order
1	2	3	4	5	1	2	3	4	5	

To see an expanded version of the assessment, go to this article at HBR.org.

How to Shape Your Culture

First you must identify culture targets. The best ones have some attributes in common: They align with the company's strategic direction; they're important to execute; and they reflect the demands of the external business environment. A good target should be both specific and achievable. For example, "We value our customers" can create ambiguity and lead to inconsistent choices regarding hiring, developing leaders, and running the company. A better version might be "We build genuine and positive relationships with customers; we serve our customers with humility; and we act as ambassadors for our rich brand heritage."

To Set a Culture Target:

Understand the current culture

Examine your culture—the company's founding and heritage, its espoused values, subcultures, leadership style, and team dynamics. (Use the preceding figure to start the conversation.)

Identify your culture's strengths and examine its impact on your organization today. Interview key stakeholders and influential members of the organization as needed.

Consider strategy and the environment

Assess current and future external conditions and strategic choices and determine which cultural styles will need to be strengthened or diminished in response.

Formulate a culture target according to which styles will support future changes.

Frame the aspiration in business realities

Translate the target into organizational change priorities. It should be framed not as a culture change initiative but in terms of real-world problems to be solved and solutions that create value.

Focus on *leadership alignment*, *organizational conversations*, and *organizational design* as the levers to guide the culture's evolution.

One Company's Experience

One large company used its search for a new director as an opportunity to bridge a problematic gap between the company's culture and the board's culture. To accomplish this, the leadership first diagnosed the two cultures along with its aspirations for the new director.

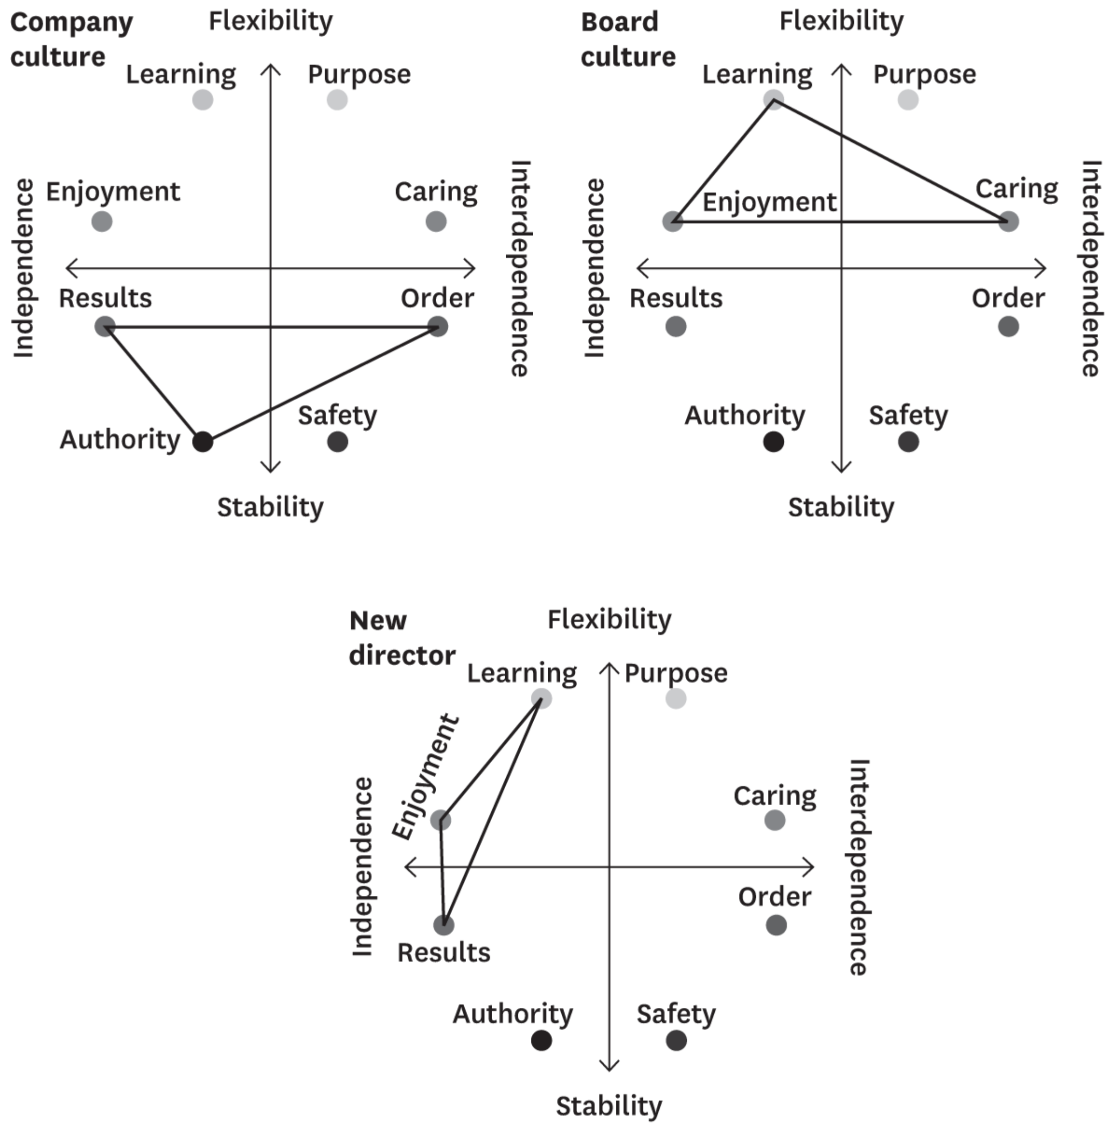

Whereas the company was highly *results* oriented and focused on *order*, discipline, and execution, the board was far more *learning* oriented, exploratory, inquisitive, and focused on *enjoyment*. A director who was *results* driven and curious would help bridge the two cultures.

Two years after an individual with the desired style was brought in, the board and the management team reported more-effective strategic planning activities and improved company performance.

Convergence Matters

When we compared employees' views on their organization's most salient cultural attributes, two types of organizations emerged: *low convergence* (employees rarely agreed on the most important cultural attributes) and *high convergence* (views were more closely aligned). In the two examples below, each dot represents one employee.

Note that in the low-convergence organization, seven of the eight cultural attributes were cited as most important, and every quadrant is represented. That means employees viewed their company in varying and often opposite ways. Some saw a *caring* organization, for example, while others saw one that emphasized *authority*.

Why is high convergence important? Because it correlates with levels of employee engagement and customer orientation. However, if the culture you have is not the one you want, high convergence will make it harder to change.

Company A: Low convergence

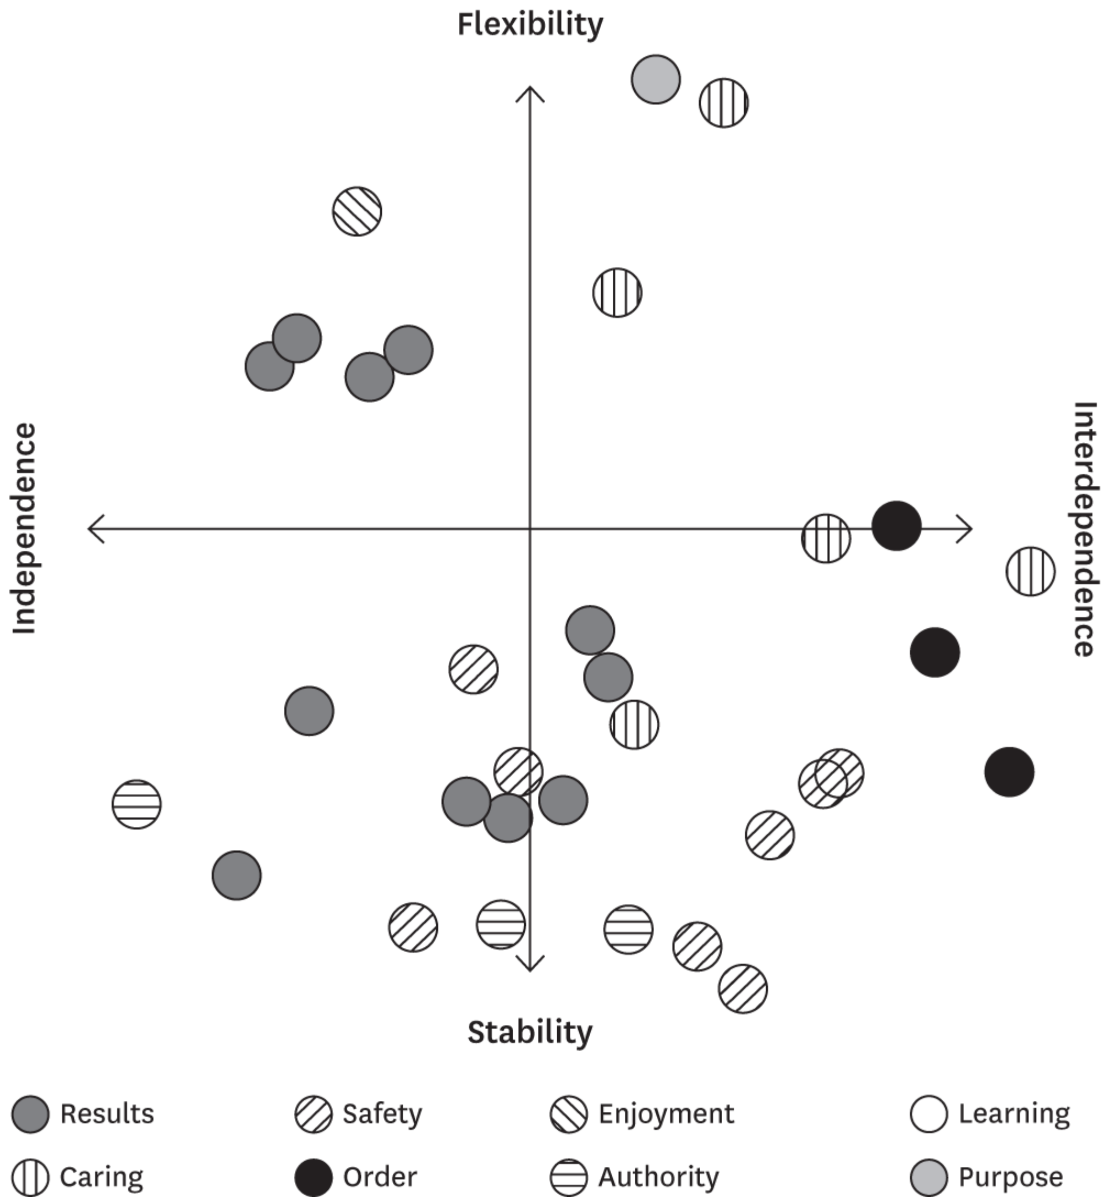

Company B: High convergence

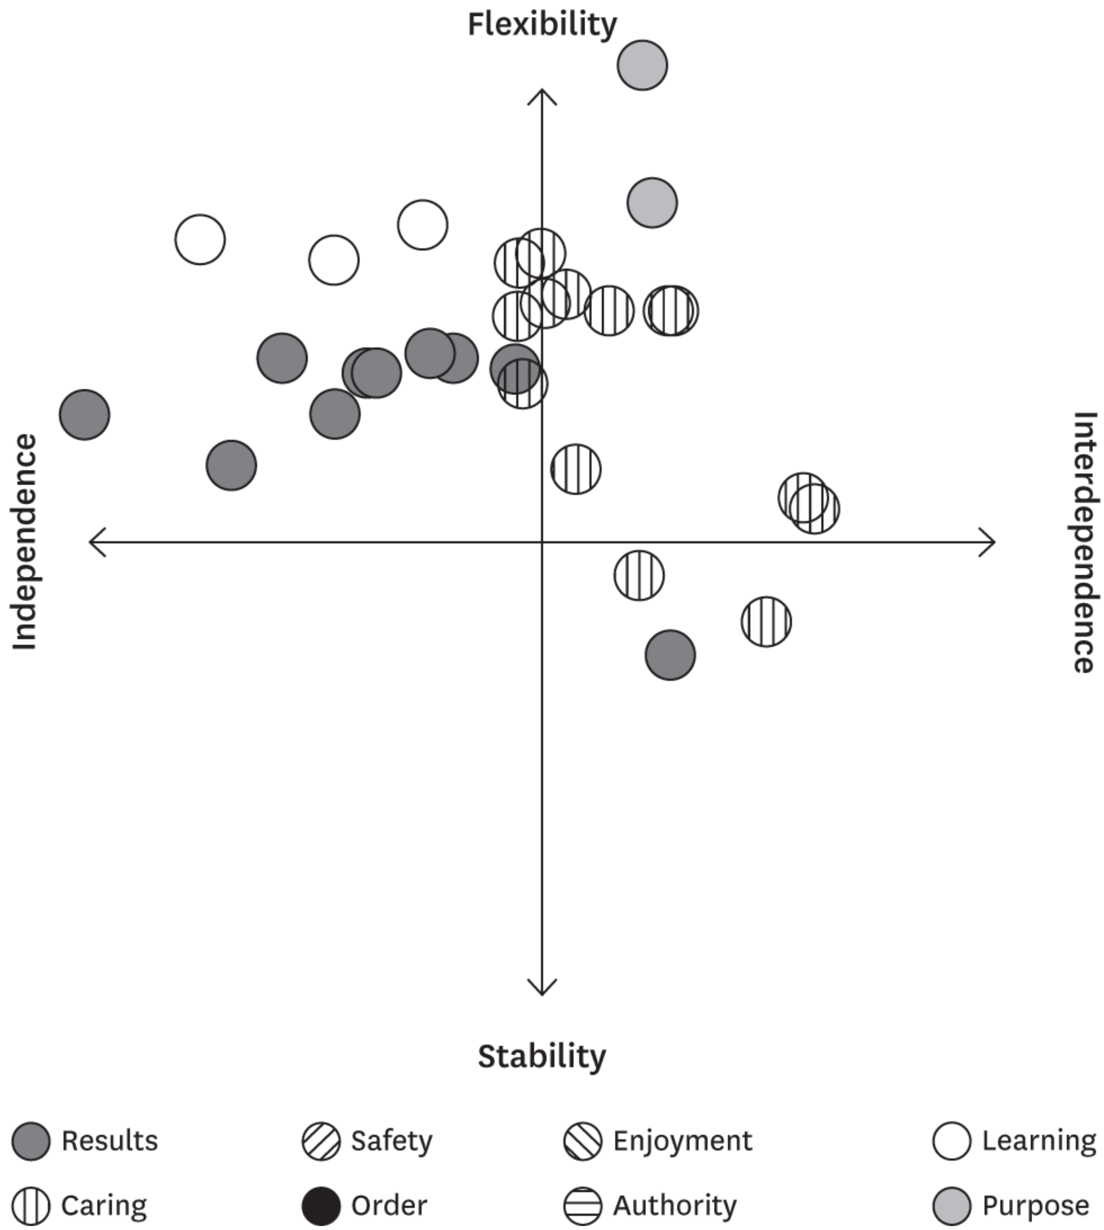

Context, Conditions, and Culture

Context matters when assessing a culture's strategic effectiveness.

Leaders must simultaneously consider culture styles and key organizational and market conditions if they want their culture to help drive performance. Region and industry are among the most germane external factors to keep in mind; critical internal considerations include alignment with strategy, leadership, and organizational design.

Region

The values of the national and regional cultures in which a company is embedded can influence patterns of behavior within the organization. (This linkage has been explored in depth by Geert Hofstede and the authors of the GLOBE study.) We find, for example, that companies operating in countries characterized by a high degree of institutional collectivism (defined as valuing equity within groups and encouraging the collective distribution of resources), such as France and Brazil, have cultures that emphasize *order* and *safety*. Companies operating in countries with low levels of uncertainty avoidance (that is, they are open to ambiguity and future uncertainty), such as the United States and Australia, place a greater emphasis on *learning*, *purpose*, and *enjoyment*. Such external influences are important considerations when working across borders or designing an appropriate organizational culture.

Industry

Varying cultural attributes may be needed to address industry-specific regulations and customer needs. A comparison of organizations across industries reveals evidence that cultures might adapt to meet the demands of industry environments.

Organizational cultures in financial services are more likely to emphasize *safety*. Given the increasingly complex regulations enacted in response to the financial crisis, careful work and risk management are more critical than ever in this industry. In contrast, nonprofits are far more purpose-driven, which can reinforce their commitment to a mission by aligning employee behavior around a common goal.

Culture styles ranked by industry

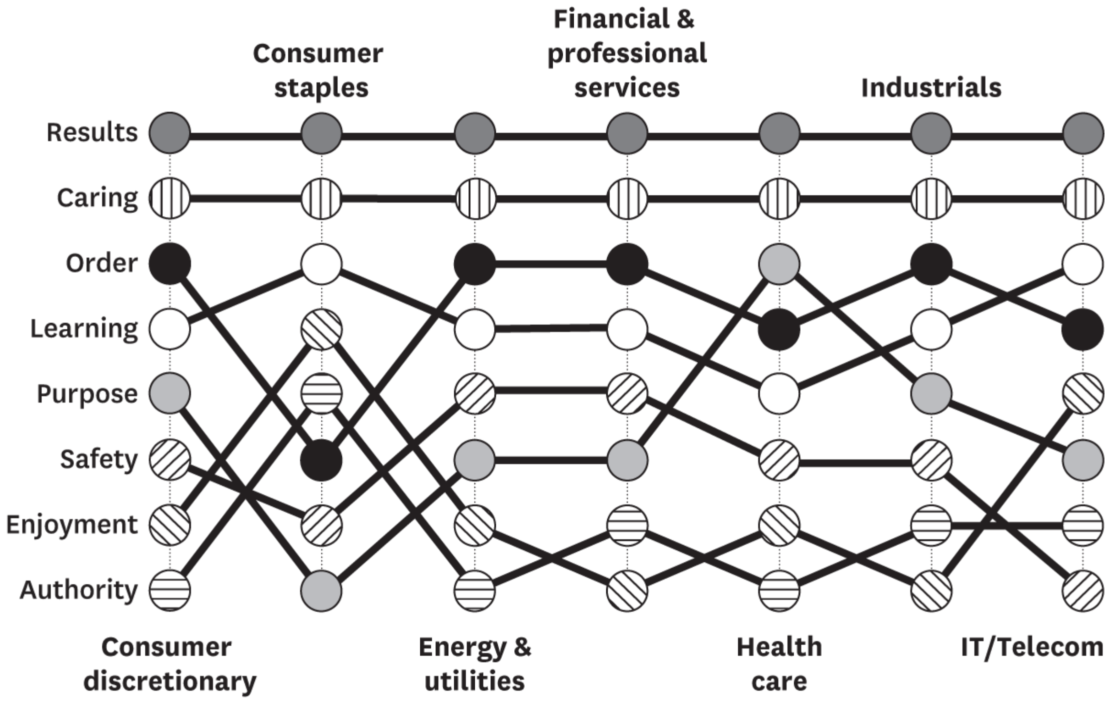

Based on an assessment of 230+ companies (industry) and a subsample of 25 companies (strategy)

Strategy

For its full benefit to be realized, a culture must support the strategic goals and plans of the business. For example, we find differences between companies that adopt a differentiation strategy and companies that pursue a cost leadership strategy. Although *results* and *caring* are key cultural characteristics at both types of companies, *enjoyment*, *learning*, and *purpose* are more suited to differentiation, whereas *order* and *authority* are more suited to cost leadership. Flexible cultures—which emphasize *enjoyment* and *learning*—can spur product innovation in companies aiming to differentiate themselves, whereas stable and predictable cultures, which emphasize *order* and *authority*, can help maintain operational efficiency to keep costs low.

Strategic considerations related to a company's life cycle are also linked to organizational culture. Companies with a strategy that seeks to stabilize or maintain their market position prioritize *learning*, whereas organizations operating with a turnaround strategy tend to prioritize *order* and *safety* in their efforts to redirect or reorganize unprofitable units.

Culture styles ranked by strategy

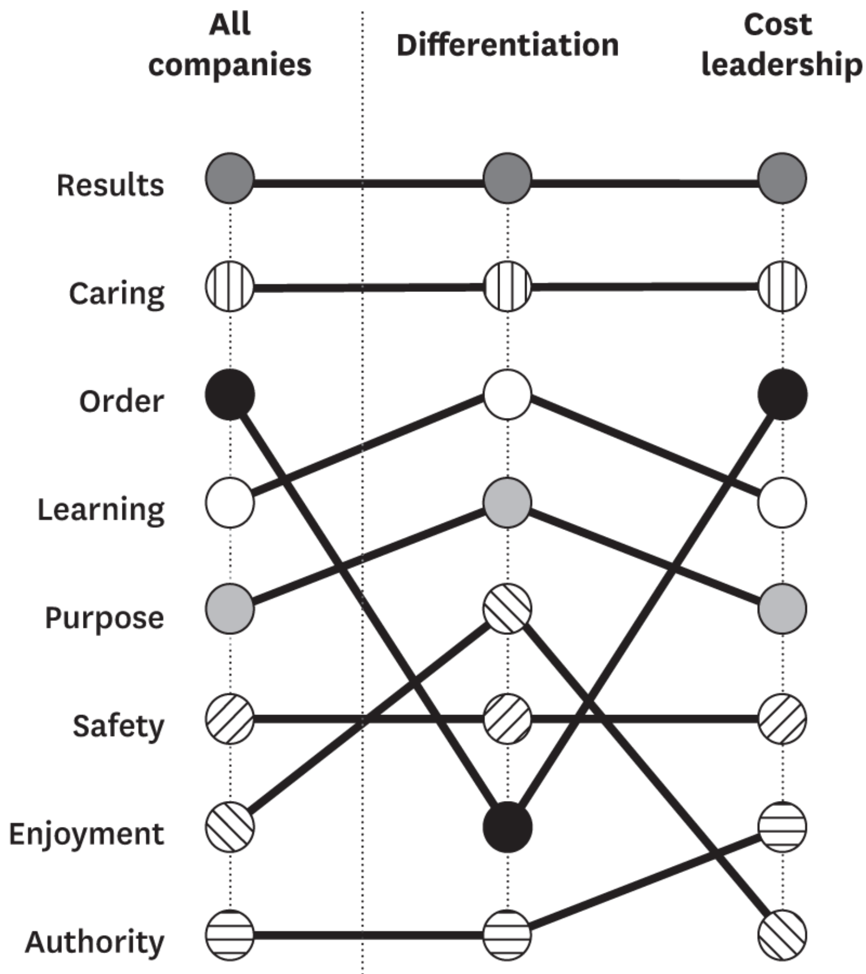

Based on an assessment of 230+ companies (industry) and a subsample of 25 companies (strategy)

Leadership

It is hard to overestimate the importance of aligning culture and leadership. The character and behaviors of a CEO and top executives can have a profound effect on culture. Conversely, culture serves to either constrain or enhance the performance of leaders. Our own data from executive recruiting activities shows that a lack of cultural fit is responsible for up to 68% of new-hire failures at the senior leadership level. For individual leaders, cultural fit is as important as capabilities and experience.

Organizational design

We see a two-way relationship between a company's culture and its particular structure. In many cases, structure and systems follow culture. For example, companies that prioritize teamwork and collaboration might design incentive systems that include shared team and company goals along with rewards that recognize collective effort. However, a long-standing organizational design choice can lead to the formation of a culture. Because the latter is far more difficult to alter, we suggest that structural changes should be aligned with the desired culture.

About the Research

We undertook a comprehensive study of organizational culture and outcomes to explore the link between them. We analyzed the cultures of more than 230 companies along with the leadership styles and values of more than 1,300 executives across a range of industries (including consumer discretionary, consumer staples, energy and utilities, financial and professional services, health care, industrials, and IT and telecommunications), regions (Africa, Asia, Europe, the Middle East, North America, Oceania, and South America), and organizational types (public, private, and nonprofit). We diagnosed those cultures using online survey responses from approximately 25,000 employees together with interviews of company managers.

Our analysis highlighted how strongly each of the eight styles defined the organizations in our study. *Results* ranked first, and *caring* second. This pattern is consistent across company types, company sizes, regions, and industries. *Order* and *learning* ranked among the third and fourth most common styles in many cultures.

Culture appears to most directly affect employee engagement and motivation, followed by customer orientation. To model its relationship to organizational outcomes, we assessed employee engagement levels for all the companies using widely accepted survey questions and arrived at customer-orientation scores with an online questionnaire. In many cases we also documented top leaders' individual styles and values.

We found that employee engagement is most strongly related to greater flexibility, in the form of *enjoyment*, *learning*, *purpose*, and *caring*. Similarly, we observed a positive relationship between customer orientation and those four styles plus *results*. These relationships, too, are surprisingly consistent across companies. We also found that engagement and customer orientation are stronger when employees are in close agreement about the culture's characteristics.

Our research was influenced by the work of countless scholars in this field, many of whom are mentioned in this article. In addition, we stand on the shoulders of giants such as David Caldwell, Jennifer Chatman, James Heskett, John Kotter, Charles O'Reilly, and many, many others who have inspired our thinking.

Originally published in January–February 2018. Reprint R1801B

The Error at the Heart of Corporate Leadership

by Joseph L. Bower and Lynn S. Paine

IN THE FALL OF 2014, the hedge fund activist and Allergan shareholder Bill Ackman became increasingly frustrated with Allergan's board of directors. In a letter to the board, he took the directors to task for their failure to do (in his words) "what you are paid $400,000 per year to do on behalf of the Company's owners." The board's alleged failure: refusing to negotiate with Valeant Pharmaceuticals about its unsolicited bid to take over Allergan—a bid that Ackman himself had helped engineer in a novel alliance between a hedge fund and a would-be acquirer. In presentations promoting the deal, Ackman praised Valeant for its shareholder-friendly capital allocation, its shareholder-aligned executive compensation, and its avoidance of risky early-stage research. Using the same approach at Allergan, he told analysts, would create significant value for its shareholders. He cited Valeant's plan to cut Allergan's research budget by 90% as "really the opportunity." Valeant CEO Mike Pearson assured analysts that "all we care about is shareholder value."

These events illustrate a way of thinking about the governance and management of companies that is now pervasive in the financial community and much of the business world. It centers on the idea that management's objective is, or should be, maximizing value for shareholders, but it addresses a wide range of topics—from performance measurement and executive compensation to shareholder rights, the role of directors, and corporate responsibility. This thought system has been embraced not only by hedge fund activists like Ackman but also by institutional investors more generally, along with many boards, managers, lawyers, academics, and even some regulators and lawmakers. Indeed, its precepts have come to be widely regarded as a model for "good governance" and for the brand of investor activism illustrated by the Allergan story.

Yet the idea that corporate managers should make maximizing shareholder value their goal—and that boards should ensure that they do—is relatively recent. It is rooted in what's known as agency theory, which was put forth by academic economists in the 1970s. At the theory's core is the assertion that shareholders own the corporation and, by virtue of their status as owners, have ultimate authority over its business and may legitimately demand that its activities be conducted in accordance with their wishes.

Attributing ownership of the corporation to shareholders sounds natural enough, but a closer look reveals that it is legally confused and, perhaps more important, involves a challenging

problem of accountability. Keep in mind that shareholders have no legal duty to protect or serve the companies whose shares they own and are shielded by the doctrine of limited liability from legal responsibility for those companies' debts and misdeeds. Moreover, they may generally buy and sell shares without restriction and are required to disclose their identities only in certain circumstances. In addition, they tend to be physically and psychologically distant from the activities of the companies they invest in. That is to say, public company shareholders have few incentives to consider, and are not generally viewed as responsible for, the effects of the actions they favor on the corporation, other parties, or society more broadly. Agency theory has yet to grapple with the implications of the accountability vacuum that results from accepting its central—and in our view, faulty—premise that shareholders own the corporation.

The effects of this omission are troubling. We are concerned that the agency-based model of governance and management is being practiced in ways that are weakening companies and—if applied even more widely, as experts predict—could be damaging to the broader economy. In particular we are concerned about the effects on corporate strategy and resource allocation. Over the past few decades the agency model has provided the rationale for a variety of changes in governance and management practices that, taken together, have increased the power and influence of certain types of shareholders over other

types and further elevated the claims of shareholders over those of other important constituencies—without establishing any corresponding responsibility or accountability on the part of shareholders who exercise that power. As a result, managers are under increasing pressure to deliver ever faster and more predictable returns and to curtail riskier investments aimed at meeting future needs and finding creative solutions to the problems facing people around the world.

Idea in Brief

The Problem

A widespread belief holds that "maximizing shareholder value" is the number one responsibility of boards and managers. But that's confused as a matter of corporate law and a poor guide for managerial behavior—and it has a huge accountability problem baked into it.

The Solution

A company's health—not its shareholders' wealth—should be the primary concern of those who manage corporations. That may sound like a small change, but it could make companies less vulnerable to damaging forms of activist investing—and make it easier for managers to focus on the long term.

Don't misunderstand: We are capitalists to the core. We believe that widespread participation in the economy through the ownership of stock in publicly traded companies is important to the social fabric, and that strong protections for shareholders are essential. But the health of the economic system depends on getting the role of shareholders right. The agency model's

extreme version of shareholder centricity is flawed in its assumptions, confused as a matter of law, and damaging in practice. A better model would recognize the critical role of shareholders but also take seriously the idea that corporations are independent entities serving multiple purposes and endowed by law with the potential to endure over time. And it would acknowledge accepted legal principles holding that directors and managers have duties to the corporation as well as to shareholders. In other words, a better model would be more company centered.

Before considering an alternative, let's take a closer look at the agency-based model.

Foundations of the Model

The ideas underlying the agency-based model can be found in Milton Friedman's well-known *New York Times Magazine* article of 1970 denouncing corporate "social responsibility" as a socialist doctrine. Friedman takes shareholders' ownership of the corporation as a given. He asserts that "the manager is the agent of the individuals who own the corporation" and, further, that the manager's primary "responsibility is to conduct the business in accordance with [the owners'] desires." He characterizes the executive as "an agent serving the interests of his principal."

These ideas were further developed in the 1976 *Journal of Financial Economics* article "Theory of the Firm," by Michael

Jensen and William Meckling, who set forth the theory's basic premises:

- Shareholders own the corporation and are "principals" with original authority to manage the corporation's business and affairs.
- Managers are delegated decision-making authority by the corporation's shareholders and are thus "agents" of the shareholders.
- As agents of the shareholders, managers are obliged to conduct the corporation's business in accordance with shareholders' desires.
- Shareholders want business to be conducted in a way that maximizes their own economic returns. (The assumption that shareholders are unanimous in this objective is implicit throughout the article.)

Jensen and Meckling do not discuss shareholders' wishes regarding the ethical standards that managers should observe in conducting the business, but Friedman offers two views in his *Times* article. First he writes that shareholders generally want managers "to make as much money as possible while conforming to the basic rules of the society, both those embodied in law and those embodied in ethical custom." Later he suggests that shareholders simply want managers to use resources and pursue profit by engaging "in open and free competition without

deception or fraud." Jensen and Meckling agree with Friedman that companies should not engage in acts of "social responsibility."

Much of the academic work on agency theory in the decades since has focused on ensuring that managers seek to maximize shareholder returns—primarily by aligning their interests with those of shareholders. These ideas have been further developed into a theory of organization whereby managers can (and should) instill concern for shareholders' interests throughout a company by properly delegating "decision rights" and creating appropriate incentives. They have also given rise to a view of boards of directors as an organizational mechanism for controlling what's known as "agency costs"—the costs to shareholders associated with delegating authority to managers. Hence the notion that a board's principal role is (or should be) monitoring management, and that boards should design executive compensation to align management's interests with those of shareholders.

The Model's Flaws

Let's look at where these ideas go astray.

1. Agency theory is at odds with corporate law: Legally, shareholders do not have the rights of "owners" of the corporation, and managers are not shareholders' "agents."

As other scholars and commentators have noted, the idea that shareholders own the corporation is at best confusing and at

worst incorrect. From a legal perspective, shareholders are beneficiaries of the corporation's activities, but they do not have "dominion" over a piece of property. Nor do they enjoy access to the corporate premises or use of the corporation's assets. What shareholders do own is their shares. That generally gives them various rights and privileges, including the right to sell their shares and to vote on certain matters, such as the election of directors, amendments to the corporate charter, and the sale of substantially all the corporation's assets.

Furthermore, under the law in Delaware—legal home to more than half the *Fortune* 500 and the benchmark for corporate law—the right to manage the business and affairs of the corporation is vested in a board of directors elected by the shareholders; the board delegates that authority to corporate managers.

Within this legal framework, managers and directors are fiduciaries rather than agents—and not just for shareholders but also for the corporation. The difference is important. Agents are obliged to carry out the wishes of a principal, whereas a fiduciary's obligation is to exercise independent judgment on behalf of a beneficiary. Put differently, an agent is an order taker, whereas a fiduciary is expected to make discretionary decisions. Legally, directors have a fiduciary duty to act in the best interests of the corporation, which is very different from simply doing the bidding of shareholders.

2. The theory is out of step with ordinary usage: Shareholders are not owners of the corporation in any traditional sense of the term,

nor do they have owners' traditional incentives to exercise care in managing it.

This observation is even truer today than when it was famously made by Adolf Berle and Gardiner Means in their landmark 1932 study *The Modern Corporation and Private Property*. Some 70% of shares in U.S.-listed companies today are held by mutual funds, pension funds, insurance companies, sovereign funds, and other institutional investors, which manage them on behalf of beneficiaries such as households, pensioners, policy holders, and governments. In many instances the beneficiaries are anonymous to the company whose shares the institutions hold. The professionals who manage these investments are typically judged and rewarded each quarter on the basis of returns from the total basket of investments managed. A consequence is high turnover in shares (seen in the exhibit "Average holding period for public company shares"), which also results from high-frequency trading by speculators.

Average holding period for public company shares

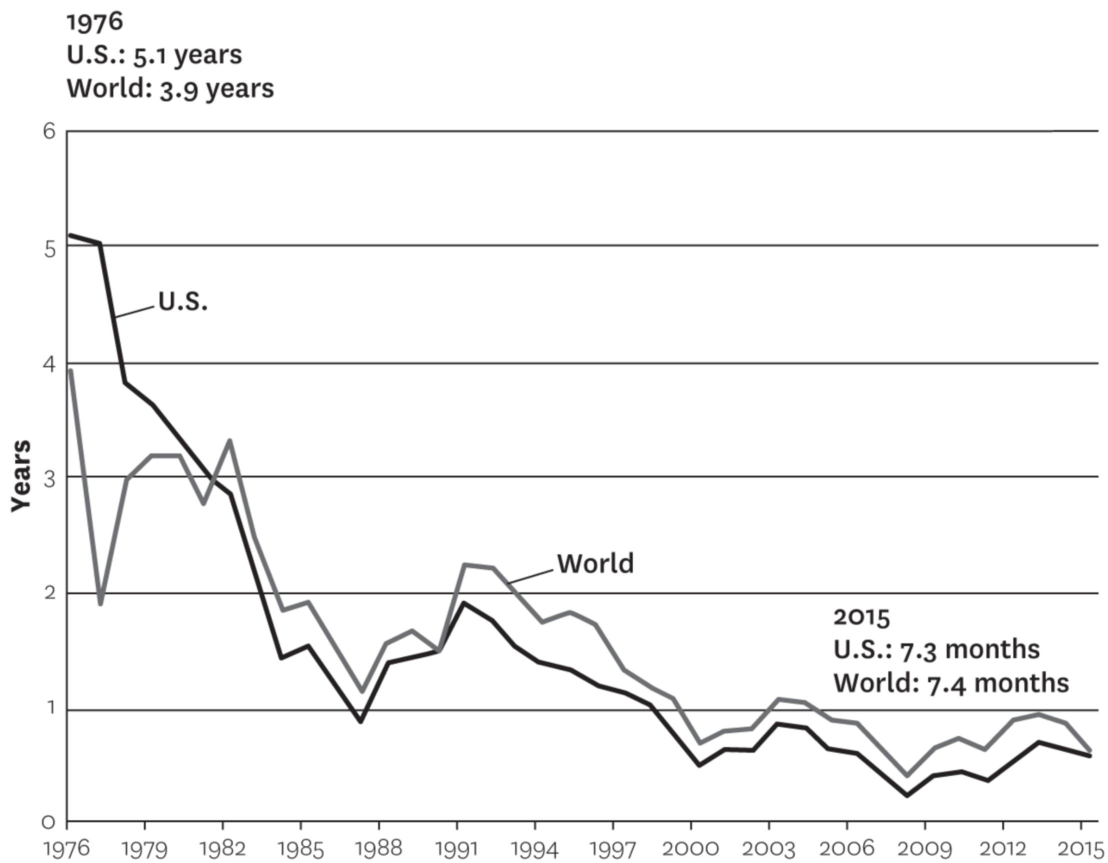

Source: The World Bank, World Federation of Exchanges Database.

The decisions of asset managers and speculators arise from expectations regarding share price over a relatively short period of time. As the economy passes through cycles, the shares of companies in entire industry sectors move in and out of favor. Although the shareholders of record at any given moment may vote on an issue brought before them, they need not know or care about the company whose shares they hold. Moreover, the fact that they can hedge or immediately sell their shares and avoid exposure to the longer-term effects of that vote makes it

difficult to regard them as proprietors of the company in any customary sense.

The anonymity afforded the shares' beneficial owners further attenuates their relationship to the companies whose shares they own. Some 85% of publicly traded shares in the United States are held in the name of an institution serving as an intermediary—the so-called street name—on behalf of itself or its customers. And of the ultimate owners of those shares, an estimated 75% have instructed their intermediaries not to divulge their identities to the issuing company.

3. The theory is rife with moral hazard: Shareholders are not accountable as owners for the company's activities, nor do they have the responsibilities that officers and directors do to protect the company's interests.

The problem with treating shareholders as proprietors is exacerbated by the absence of another traditional feature of ownership: responsibility for the property owned and accountability—even legal liability, in some cases—for injuries to third parties resulting from how that property is used. Shareholders bear no such responsibility. Under the doctrine of limited liability, they cannot be held personally liable for the corporation's debts or for corporate acts and omissions that result in injury to others.

With a few exceptions, shareholders are entitled to act entirely in their own interest within the bounds of the securities laws. Unlike directors, who are expected to refrain from self-dealing,

they are free to act on both sides of a transaction in which they have an interest. Consider the contest between Allergan and Valeant. A member of Allergan's board who held shares in Valeant would have been expected to refrain from voting on the deal or promoting Valeant's bid. But Allergan shareholders with a stake in both companies were free to buy, sell, and vote as they saw fit, with no obligation to act in the best interests of either company. Institutional investors holding shares in thousands of companies regularly act on deals in which they have significant interests on both sides.

In a well-ordered economy, rights and responsibilities go together. Giving shareholders the rights of ownership while exempting them from the responsibilities opens the door to opportunism, overreach, and misuse of corporate assets. The risk is less worrying when shareholders do not seek to influence major corporate decisions, but it is acute when they do. The problem is clearest when temporary holders of large blocks of shares intervene to reconstitute a company's board, change its management, or restructure its finances in an effort to drive up its share price, only to sell out and move on to another target without ever having to answer for their intervention's impact on the company or other parties.

4. The theory's doctrine of alignment spreads moral hazard throughout a company and narrows management's field of vision.
Just as freedom from accountability has a tendency to make shareholders indifferent to broader and longer-term

considerations, so agency theory's recommended alignment between managers' interests and those of shareholders can skew the perspective of the entire organization. When the interests of successive layers of management are "aligned" in this manner, the corporation may become so biased toward the narrow interests of its current shareholders that it fails to meet the requirements of its customers or other constituencies. In extreme cases it may tilt so far that it can no longer function effectively. The story of Enron's collapse reveals how thoroughly the body of a company can be infected.

The notion that managing for the good of the company is the same as managing for the good of the stock is best understood as a theoretical conceit necessitated by the mathematical models that many economists favor. In practical terms there is (or can be) a stark difference. Once Allergan's management shifted its focus from sustaining long-term growth to getting the company's stock price to $180 a share—the target at which institutional investors were willing to hold their shares—its priorities changed accordingly. Research was cut, investments were eliminated, and employees were dismissed.

5. The theory's assumption of shareholder uniformity is contrary to fact: Shareholders do not all have the same objectives and cannot be treated as a single "owner."

Agency theory assumes that all shareholders want the company to be run in a way that maximizes their own economic return. This simplifying assumption is useful for certain purposes, but it

masks important differences. Shareholders have differing investment objectives, attitudes toward risk, and time horizons. Pension funds may seek current income and preservation of capital. Endowments may seek long-term growth. Young investors may accept considerably more risk than their elders will tolerate. Proxy voting records indicate that shareholders are divided on many of the resolutions put before them. They may also view strategic opportunities differently. In the months after Valeant announced its bid, Allergan officials met with a broad swath of institutional investors. According to Allergan's lead independent director, Michael Gallagher, "The diversity of opinion was as wide as could possibly be"—from those who opposed the deal and absolutely did not want Valeant shares (the offer included both stock and cash) to those who saw it as the opportunity of a lifetime and could not understand why Allergan did not sit down with Valeant immediately.

The Agency-Based Model in Practice

Despite these problems, agency theory has attracted a wide following. Its tenets have provided the intellectual rationale for a variety of changes in practice that, taken together, have enhanced the power of shareholders and given rise to a model of governance and management that is unrelenting in its shareholder centricity. Here are just a few of the arenas in which the theory's influence can be seen:

Executive compensation

Agency theory ideas were instrumental in the shift from a largely cash-based system to one that relies predominantly on equity. Proponents of the shift argued that equity-based pay would better align the interests of executives with those of shareholders. The same argument was used to garner support for linking pay more closely to stock performance and for tax incentives to encourage such "pay for performance" arrangements. Following this logic, Congress adopted legislation in 1992 making executive pay above $1 million deductible only if it is "performance based." Today some 62% of executive pay is in the form of equity, compared with 19% in 1980.

Disclosure of executive pay

Agency theory's definition of performance and its doctrine of alignment undergird rules proposed by the SEC in 2015 requiring companies to expand the information on executive pay and shareholder returns provided in their annual proxy statements. The proposed rules call for companies to report their annual total shareholder return (TSR) over time, along with annual TSR figures for their peer group, and to describe the relationships between their TSR and their executive compensation and between their TSR and the TSR of their peers.

Shareholders' rights

The idea that shareholders are owners has been central to the push to give them more say in the nomination and election of

directors and to make it easier for them to call a special meeting, act by written consent, or remove a director. Data from FactSet and other sources indicates that the proportion of S&P 500 companies with majority voting for directors increased from about 16% in 2006 to 88% in 2015; the proportion with special meeting provisions rose from 41% in 2002 to 61% in 2015; and the proportion giving shareholders proxy access rights increased from less than half a percent in 2013 to some 39% by mid-2016.

The power of boards

Agency thinking has also propelled efforts to eliminate staggered boards in favor of annual election for all directors and to eliminate “poison pills” that would enable boards to slow down or prevent “owners” from voting on a premium offer for the company. From 2002 to 2015, the share of S&P 500 companies with staggered boards dropped from 61% to 10%, and the share with a standing poison pill fell from 60% to 4%. (Companies without a standing pill may still adopt a pill in response to an unsolicited offer—as was done by the Allergan board in response to Valeant’s bid.)

Management attitudes

Agency theory’s conception of management responsibility has been widely adopted. In 1997 the Business Roundtable issued a statement declaring that “the paramount duty of management and of boards of directors is to the corporation’s stockholders” and that “the principal objective of a business enterprise is to

generate economic returns to its owners." Issued in response to pressure from institutional investors, the statement in effect revised the Roundtable's earlier position that "the shareholder must receive a good return but the legitimate concerns of other constituencies also must have the appropriate attention." Various studies suggest ways in which managers have become more responsive to shareholders. Research indicates, for instance, that companies with majority (rather than plurality) voting for directors are more apt to adopt shareholder proposals that garner majority support, and that many chief financial officers are willing to forgo investments in projects expected to be profitable in the longer term in order to meet analysts' quarterly earnings estimates. According to surveys by the Aspen Institute, many business school graduates regard maximizing shareholder value as their top responsibility.

Investor behavior

Agency theory ideas have facilitated a rise in investor activism and legitimized the playbook of hedge funds that mobilize capital for the express purpose of buying company shares and using their position as "owners" to effect changes aimed at creating shareholder value. (The sidebar "The Activist's Playbook" illustrates how agency theory ideas have been put into practice.) These investors are intervening more frequently and reshaping how companies allocate resources. In the process they are reshaping the strategic context in which all companies and their boards make decisions.

The Activist's Playbook

For an understanding of the agency-based model in practice, there is no better place to look than an activist campaign. As a first step, the activist acquires shares in the targeted company—typically somewhere between 5% and 10%, but sometimes less than 1%. Shares in hand, he then claims the right to issue directives. (To leverage that power, he will often alert other hedge funds to his actions.) The language of ownership typically plays a prominent role. For example, in 2014, to advance a takeover of Allergan by Valeant Pharmaceuticals, Bill Ackman, of Pershing Square Capital Management, attacked Allergan's board for failing to do what the directors were paid to do "on behalf of the Company's owners." The activist may challenge the board's professionalism by appealing to agency theory norms of directorship. In one letter to the Allergan board, Ackman declared: "Your actions have wasted corporate resources, delayed enormous potential value creation for shareholders, and are professionally and personally embarrassing for you."

Although campaigns differ in their particulars, the activist's playbook for increasing shareholder value is fairly standard. As our colleagues Ian Gow and Suraj Srinivasan (with others) have documented in their study of nearly 800 campaigns at U.S. companies from 2004 to 2012, activists tend to focus on capital structure, strategy, and governance. They typically call for some combination of cutting costs, adding debt, buying back shares, issuing special dividends, spinning off businesses, reconstituting the board, replacing the CEO, changing the strategy, and selling the company or its main asset. Tax reduction is another element of many activist programs.

An activist whose demands go unheeded may initiate a proxy fight in an attempt to replace incumbent board members with directors more willing to do the activist's bidding. In a few instances, activists have even offered their chosen nominees special bonuses to stand for

election or additional incentives for increasing shareholder value in their role as directors.

By most indications, hedge fund activists have been quite successful in effecting the changes they've sought. As reported by the industry, more companies are being targeted—473 worldwide in the first half of 2016 (including 306 in the United States), up from 136 worldwide in all of 2010—and activists' demands are frequently being met. In the United States in 2015, 69% of demands were at least partially satisfied, the highest proportion since 2010. Activists are also gaining clout in the boardroom, where they won 397 seats at U.S. companies in 2014 and 2015. Although activist hedge funds saw outflows of some $7.4 billion in the first three quarters of 2016, assets under management were estimated at more than $116 billion in late 2016, up from $2.7 billion in 2000.

Taken individually, a change such as majority voting for directors may have merit. As a group, however, these changes have helped create an environment in which managers are under increasing pressure to deliver short-term financial results, and boards are being urged to "think like activists."

Implications for Companies

To appreciate the strategic implications of a typical activist program, it is instructive to use a tool developed in the 1960s by the Boston Consulting Group to guide the resource-allocation process. Called the growth share matrix, the tool helped managers see their company as a portfolio of businesses with differing characteristics. One group of businesses might be mature and require investment only for purposes of modest

expansion and incremental improvement. Assuming they have strong market share relative to their nearest competitors, those businesses are likely to be profitable and generate cash. Another group might also have leading positions but be in fast-growing markets; they, too, are profitable, but they require heavy investment to maintain or improve market share. A third group might have weak competitive positions in mature markets; these businesses require cash for survival but have no prospects for growth or increased profits. A final group might be in rapidly growing new markets where several companies are competitive and prospects are bright but risky.

The developers of the matrix called these four groups cash cows, stars, dogs, and bright prospects, respectively. The segmentation was meant to ensure that cash cows were maintained, stars fully funded, dogs pruned, and a limited number of bright prospects chosen for their longer-term potential to become stars. (See the exhibit "The growth share matrix.") When companies don't manage a portfolio in this holistic fashion, funds tend to get spread evenly across businesses on the basis of individual projects' forecasted returns.

The growth share matrix

BCG's growth share matrix enables companies to manage a portfolio of businesses: "cash cows," mature businesses that throw off cash; fast-growing "stars"; businesses with a weak position and few prospects for growth ("dogs"); and risky but big-upside businesses in fast-growing markets ("bright prospects").

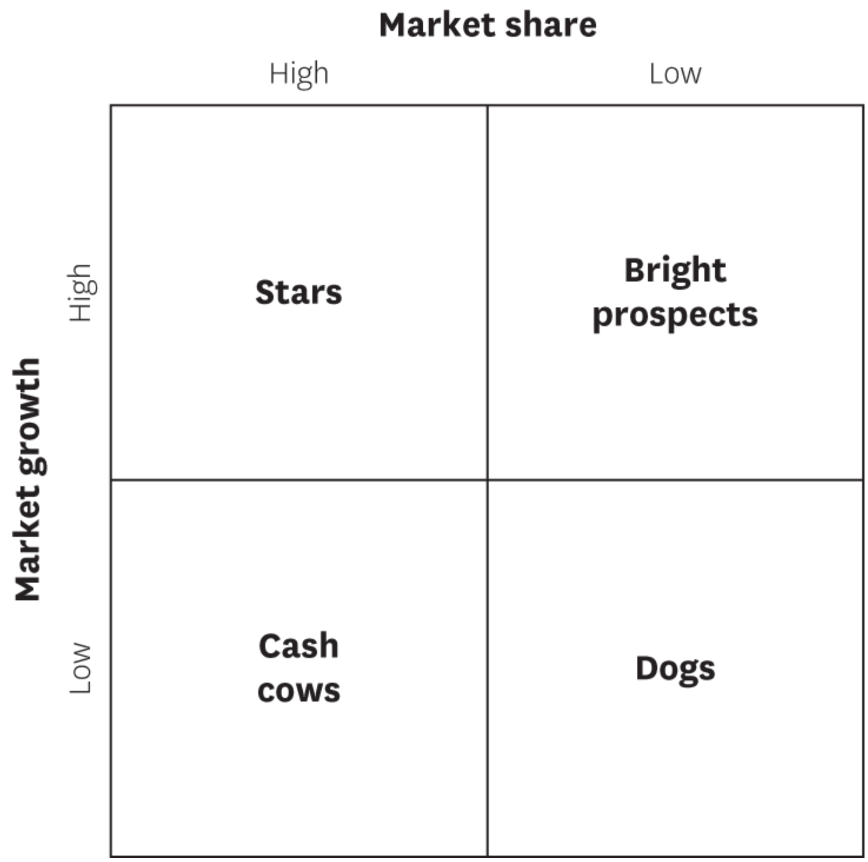

Source: Boston Consulting Group.

It's a simple tool—but using it well is not simple at all. Managing a cash cow so that it remains healthy, nurturing star businesses in the face of emerging competition, fixing or divesting unpromising businesses, and selecting one or two bright prospects to grow—all this takes talented executives who can function effectively as a team. Companies that succeed in managing this ongoing resource-allocation challenge can grow and reinvent themselves continually over time.

The growth share matrix illuminates the strategic choices managers face as they seek to create value indefinitely into the

future. It's also useful for showing how to drive up a company's share price in the short term. Suppose a corporation were to sell off the dogs, defund the bright prospects, and cut expenses such as marketing and R&D from the stars. That's a recipe for dramatically increased earnings, which would, in turn, drive up the share price. But the corporation might lose bright prospects that could have been developed into the stars and cash cows of the future.

The activist investor Nelson Peltz's 2014 proposal for DuPont provides an example of this idea. At the core of his three-year plan for increasing returns to shareholders was splitting the company into three autonomous businesses and eliminating its central research function. One of the new companies, "GrowthCo," was to consist of DuPont's agriculture, nutrition and health, and industrial biosciences businesses. A second, "CyclicalCo/CashCo," was to include the low-growth but highly cash-generative performance materials, safety, and electronics businesses. The third was the performance chemicals unit, Chemours, which DuPont had already decided to spin off. In growth-share-matrix terms, Peltz's plan was, in essence, to break up DuPont into a cash cow, a star, and a dog—and to eliminate some number of the bright prospects that might have been developed from innovations produced by centralized research. Peltz also proposed cutting other "excess" costs, adding debt, adopting a more shareholder-friendly policy for distributing cash from CyclicalCo/CashCo, prioritizing high returns on invested

capital for initiatives at GrowthCo, and introducing more shareholder-friendly governance, including tighter alignment between executive compensation and returns to shareholders. The plan would effectively dismantle DuPont and cap its future in return for an anticipated doubling in share price.

Value Creation or Value Transfer?

The question of whether shareholders benefit from such activism beyond an initial bump in stock price is likely to remain unresolved, given the methodological problems plaguing studies on the subject. No doubt in some cases activists have played a useful role in waking up a sleepy board or driving a long-overdue change in strategy or management. However, it is important to note that much of what activists call value creation is more accurately described as value transfer. When cash is paid out to shareholders rather than used to fund research, launch new ventures, or grow existing businesses, value has not been created. Nothing has been created. Rather, cash that would have been invested to generate future returns is simply being paid out to current shareholders. The lag time between when such decisions are taken and when their effect on earnings is evident exceeds the time frames of standard financial models, so the potential for damage to the company and future shareholders, not to mention society more broadly, can easily go unnoticed.

Given how long it takes to see the fruits of any significant research effort (Apple's latest iPhone chip was eight years in the

making), the risk to research and innovation from activists who force deep cuts to drive up the share price and then sell out before the pipeline dries up is obvious. It doesn't help that financial models and capital markets are notoriously poor at valuing innovation. After Allergan was put into play by the offer from Valeant and Ackman's Pershing Square Capital Management, the company's share price rose by 30% as other hedge funds bought the stock. Some institutions sold to reap the immediate gain, and Allergan's management was soon facing pressure from the remaining institutions to accelerate cash flow and "bring earnings forward." In an attempt to hold on to those shareholders, the company made deeper cuts in the workforce than previously planned and curtailed early-stage research programs. Academic studies have found that a significant proportion of hedge fund interventions involve large increases in leverage and large decreases in investment, particularly in research and development.

The activists' claim of value creation is further clouded by indications that some of the value purportedly created for shareholders is actually value transferred from other parties or from the general public. Large-sample research on this question is limited, but one study suggests that the positive abnormal returns associated with the announcement of a hedge fund intervention are, in part, a transfer of wealth from workers to shareholders. The study found that workers' hours decreased and their wages stagnated in the three years after an

intervention. Other studies have found that some of the gains for shareholders come at the expense of bondholders. Still other academic work links aggressive pay-for-stock-performance arrangements to various misdeeds involving harm to consumers, damage to the environment, and irregularities in accounting and financial reporting.

We are not aware of any studies that examine the total impact of hedge fund interventions on all stakeholders or society at large. Still, it appears self-evident that shareholders' gains are sometimes simply transfers from the public purse, such as when management improves earnings by shifting a company's tax domicile to a lower-tax jurisdiction—a move often favored by activists, and one of Valeant's proposals for Allergan. Similarly, budget cuts that eliminate exploratory research aimed at addressing some of society's most vexing challenges may enhance current earnings but at a cost to society as well as to the company's prospects for the future.

Hedge fund activism points to some of the risks inherent in giving too much power to unaccountable "owners." As our analysis of agency theory's premises suggests, the problem of moral hazard is real—and the consequences are serious. Yet practitioners continue to embrace the theory's doctrines; regulators continue to embed them in policy; boards and managers are under increasing pressure to deliver short-term returns; and legal experts forecast that the trend toward greater shareholder empowerment will persist. To us, the prospect that

public companies will be run even more strictly according to the agency-based model is alarming. Rigid adherence to the model by companies uniformly across the economy could easily result in even more pressure for current earnings, less investment in R&D and in people, fewer transformational strategies and innovative business models, and further wealth flowing to sophisticated investors at the expense of ordinary investors and everyone else.

Toward a Company-Centered Model

A better model, we submit, would have at its core the health of the enterprise rather than near-term returns to its shareholders. Such a model would start by recognizing that corporations are independent entities endowed by law with the potential for indefinite life. With the right leadership, they can be managed to serve markets and society over long periods of time. Agency theory largely ignores these distinctive and socially valuable features of the corporation, and the associated challenges of managing for the long term, on the grounds that corporations are "legal fictions." In their seminal 1976 article, Jensen and Meckling warn against "falling into the trap" of asking what a company's objective should be or whether the company has a social responsibility. Such questions, they argue, mistakenly imply that a corporation is an "individual" rather than merely a convenient legal construct. In a similar vein, Friedman asserts

that it cannot have responsibilities because it is an “artificial person.”

In fact, of course, corporations *are* legal constructs, but that in no way makes them artificial. They are economic and social organisms whose creation is authorized by governments to accomplish objectives that cannot be achieved by more-limited organizational forms such as partnerships and proprietorships. Their nearly 400-year history of development speaks to the important role they play in society. Originally a corporation’s objectives were set in its charter—build and operate a canal, for example—but eventually the form became generic so that corporations could be used to accomplish a wide variety of objectives chosen by their management and governing bodies. As their scale and scope grew, so did their power. The choices made by corporate decision makers today can transform societies and touch the lives of millions, if not billions, of people across the globe.

The model we envision would acknowledge the realities of managing these organizations over time and would be responsive to the needs of all shareholders—not just those who are most vocal at a given moment. Here we offer eight propositions that together provide a radically different and, we believe, more realistic foundation for corporate governance and shareholder engagement.

1. Corporations are complex organizations whose effective functioning depends on talented leaders and managers.

The success of a leader has more to do with intrinsic motivation, skills, capabilities, and character than with whether his or her pay is tied to shareholder returns. If leaders are poorly equipped for the job, giving them more "skin in the game" will not improve the situation and may even make it worse. (Part of the problem with equity-based pay is that it conflates executive skill and luck.) The challenges of corporate leadership—crafting strategy, building a strong organization, developing and motivating talented executives, and allocating resources among the corporation's various businesses for present and future returns—are significant. In focusing on incentives as the key to ensuring effective leadership, agency theory diminishes these challenges and the importance of developing individuals who can meet them.

2. Corporations can prosper over the long term only if they're able to learn, adapt, and regularly transform themselves.

In some industries today, companies may need reinvention every five years to keep up with changes in markets, competition, or technology. Changes of this sort, already difficult, are made more so by the idea that management is about assigning individuals fixed decision rights, giving them clear goals, offering them incentives to achieve those goals, and then paying them (or not) depending on whether the goals are met. This approach presupposes a degree of predictability, hierarchy, and task independence that is rare in today's organizations. Most tasks involve cooperation across

organizational lines, making it difficult to establish clear links between individual contributions and specific outcomes.

3. Corporations perform many functions in society.

One of them is providing investment opportunities and generating wealth, but corporations also produce goods and services, provide employment, develop technologies, pay taxes, and make other contributions to the communities in which they operate. Singling out any one of these as “the purpose of the corporation” may say more about the commentator than about the corporation. Agency economists, it seems, gravitate toward maximizing shareholder wealth as the central purpose. Marketers tend to favor serving customers. Engineers lean toward innovation and excellence in product performance. From a societal perspective, the most important feature of the corporation may be that it performs all these functions simultaneously over time. As a historical matter, the original purpose of the corporation—reflected in debates about limited liability and general incorporation statutes—was to facilitate economic growth by enabling projects that required large-scale, long-term investment.

4. Corporations have differing objectives and differing strategies for achieving them.

The purpose of the (generic) corporation from a societal perspective is not the same as the purpose of a (particular) corporation as seen by its founders, managers, or governing

authorities. Just as the purposes and strategies of individual companies vary widely, so must their performance measures. Moreover, companies' strategies are almost always in transition as markets change. An overemphasis on TSR for assessing and comparing corporate performance can distort the allocation of resources and undermine a company's ability to deliver on its chosen strategy.

5. Corporations must create value for multiple constituencies.
In a free market system, companies succeed only if customers want their products, employees want to work for them, suppliers want them as partners, shareholders want to buy their stock, and communities want their presence. Figuring out how to maintain these relationships and deciding when trade-offs are necessary among the interests of these various groups are central challenges of corporate leadership. Agency theory's implied decision rule—that managers should always maximize value for shareholders—oversimplifies this challenge and leads eventually to systematic underinvestment in other important relationships.

6. Corporations must have ethical standards to guide interactions with all their constituencies, including shareholders and society at large.
Adherence to these standards, which go beyond forbearance from fraud and collusion, is essential for earning the trust companies need to function effectively over time. Agency theory's ambivalence regarding corporate ethics can set

companies up for destructive and even criminal behavior—which generates a need for the costly regulations that agency theory proponents are quick to decry.

7. Corporations are embedded in a political and socioeconomic system whose health is vital to their sustainability.

Elsewhere we have written about the damaging and often self-destructive consequences of companies' indifference to negative externalities produced by their activities. We have also found that societal and systemwide problems can be a source of both risk and opportunity for companies. Consider Ecomagination, the business GE built around environmental challenges, or China Mobile's rural communications strategy, which helped narrow the digital divide between China's urban and rural populations and fueled the company's growth for nearly half a decade. Agency theory's insistence that corporations (because they are legal fictions) cannot have social responsibilities and that societal problems are beyond the purview of business (and should be left to governments) results in a narrowness of vision that prevents corporate leaders from seeing, let alone acting on, many risks and opportunities.

8. The interests of the corporation are distinct from the interests of any particular shareholder or constituency group.

As early as 1610, the directors of the Dutch East India Company recognized that shareholders with a 10-year time horizon would be unenthusiastic about the company's investing resources in

longer-term projects that were likely to pay off only in the second of two 10-year periods allowed by the original charter. The solution, suggested one official, was to focus not on the initial 10-year investors but on the strategic goals of the enterprise, which in this case meant investing in those longer-term projects to maintain the company's position in Asia. The notion that all shareholders have the same interests and that those interests are the same as the corporation's masks such fundamental differences. It also provides intellectual cover for powerful shareholders who seek to divert the corporation to their own purposes while claiming to act on behalf of all shareholders.

These propositions underscore the need for an approach to governance that takes the corporation seriously as an institution in society and centers on the sustained performance of the enterprise. They also point to a stronger role for boards and a system of accountability for boards and executives that includes but is broader than accountability to shareholders. In the model implied by these propositions, boards and business leaders would take a fundamentally different approach to such basic tasks as strategy development, resource allocation, performance evaluation, and shareholder engagement. For instance, managers would be expected to take a longer view in formulating strategy and allocating resources.

The new model has yet to be fully developed, but its conceptual foundations can be outlined. As shown in the exhibit

“Contrasting approaches to corporate governance,” the company-centered model we envision tracks basic corporate law in holding that a corporation is an independent entity, that management’s authority comes from the corporation’s governing body and ultimately from the law, and that managers are fiduciaries (rather than agents) and are thus obliged to act in the best interests of the corporation and its shareholders (which is not the same as carrying out the wishes of even a majority of shareholders). This model recognizes the diversity of shareholders’ goals and the varied roles played by corporations in society. We believe that it aligns better than the agency-based model does with the realities of managing a corporation for success over time and is thus more consistent with corporations’ original purpose and unique potential as vehicles for projects involving large-scale, long-term investment.

Contrasting approaches to corporate governance

	Shareholder centered	Company centered
Theory	***Agency theory***	***Entity theory***
Conception of the corporation	Legal fiction; nexus of contracts; pool of capital	Legal entity; social and economic organism; purposeful organization
Origins of the corporation	Private agreement among property owners to pool and increase capital	Created by lawmakers to encourage investment in long-term, large-scale projects needed by society
Functions of the corporation	Maximize wealth for shareholders	Provide goods and services; provide employment; create opportunities for investment; drive innovation
Purpose of specific corporations	Maximize shareholder value	Business purpose set by the particular company's board
Responsibilities to society	None (fictional entities can't have responsibilities)	Fulfill business purpose and act as a good corporate citizen
Ethical standards	Unclear: whatever shareholders want, or obey law and avoid fraud or collusion	Obey law and follow generally accepted ethical standards
Role of shareholders	Principals/owners of the corporation with authority over its business	Owners of shares; suppliers of capital with defined rights and responsibilities
Nature of shareholders	Undifferentiated, self-interested wealth maximizers	Diverse, with differing objectives, incentives, time horizons, and preferences
Role of directors	Shareholders' agents, delegates, or representatives	Fiduciaries for the corporation and its shareholders
Role of management	Shareholders' agents	Leaders of the organization; fiduciaries for the corporation and its shareholders
Management's objective	Maximize returns to shareholders	Sustain performance of the enterprise
Management's time frame	Present/near term (theory assumes the current share price captures all available knowledge about the company's future)	Established by the board; potentially indefinite, requiring attention to near, medium, and long term
Management performance metrics	Single: returns to shareholders	Multiple: returns to shareholders; company value; achievement of strategic goals; quality of goods and services; employee well-being
Strength	Simple structure permits clear economic argument	Consistent with law, history, and the realities facing managers
Weakness	Principles do not accord with law or good management; shareholders have power without accountability	Principles describe complex relationships and responsibilities; success is difficult to assess

---

The practical implications of company-centered governance are far-reaching. In boardrooms adopting this approach, we would expect to see some or all of these features:

- greater likelihood of a staggered board to facilitate continuity and the transfer of institutional knowledge

- more board-level attention to succession planning and leadership development
- more board time devoted to strategies for the company's continuing growth and renewal
- closer links between executive compensation and achieving the company's strategic goals
- more attention to risk analysis and political and environmental uncertainty
- a strategic (rather than narrowly financial) approach to resource allocation
- a stronger focus on investments in new capabilities and innovation
- more-conservative use of leverage as a cushion against market volatility
- concern with corporate citizenship and ethical issues that goes beyond legal compliance

A company-centered model of governance would not relieve corporations of the need to provide a return over time that reflected the cost of capital. But they would be open to a wider range of strategic positions and time horizons and would more easily attract investors who shared their goals. Speculators will always seek to exploit changes in share price—but it's not inevitable that they will color all corporate governance. It's just that agency theory, in combination with other doctrines of

modern economics, has erased the distinctions among investors and converted all of us into speculators.

If our model were accepted, speculators would have less opportunity to profit by transforming long-term players into sources of higher earnings and share prices in the short term. The legitimizing argument for attacks by unaccountable parties with opaque holdings would lose its force. We can even imagine a new breed of investors and asset managers who would focus explicitly on long-term investing. They might develop new valuation models that take a broader view of companies' prospects or make a specialty of valuing the hard-to-value innovations and intangibles—and also the costly externalities—that are often ignored in today's models. They might want to hold shares in companies that promise a solid and continuing return and that behave as decent corporate citizens. Proxy advisers might emerge to serve such investors.

We would also expect to find more support for measures to enhance shareholders' accountability. For instance, activist shareholders seeking significant influence or control might be treated as fiduciaries for the corporation or restricted in their ability to sell or hedge the value of their shares. Regulators might be inclined to call for greater transparency regarding the beneficial ownership of shares. In particular, activist funds might be required to disclose the identities of their investors and to provide additional information about the nature of their own governance. Regulators might close the 10-day window currently

afforded between the time a hedge fund acquires a disclosable stake and the time the holding must actually be disclosed. To date, efforts to close the window have met resistance from agency theory proponents who argue that it is needed to give hedge funds sufficient incentive to engage in costly efforts to dislodge poorly performing managers.

The time has come to challenge the agency-based model of corporate governance. Its mantra of maximizing shareholder value is distracting companies and their leaders from the innovation, strategic renewal, and investment in the future that require their attention. History has shown that with enlightened management and sensible regulation, companies can play a useful role in helping society adapt to constant change. But that can happen only if directors and managers have sufficient discretion to take a longer, broader view of the company and its business. As long as they face the prospect of a surprise attack by unaccountable "owners," today's business leaders have little choice but to focus on the here and now.

Further Reading

Below are some of the books and articles that examine themes touched on in this article.

- **Capitalism at Risk: Rethinking the Role of Business,** Joseph L. Bower, Herman B. Leonard, and Lynn S. Paine, Harvard Business

Review Press, 2011

- **Firm Commitment: Why the Corporation Is Failing Us and How to Restore Trust in It,** Colin Mayer, Oxford University Press, 2013
- **Fixing the Game: Bubbles, Crashes, and What Capitalism Can Learn from the NFL,** Roger L. Martin, Harvard Business Review Press, 2011
- **The Shareholder Value Myth: How Putting Shareholders First Harms Investors, Corporations, and the Public,** Lynn Stout, Berrett-Koehler, 2012
- **"Focusing Capital on the Long Term,"** Dominic Barton and Mark Wiseman, HBR, January–February 2014
- **"A Global Leader's Guide to Managing Business Conduct,"** Lynn S. Paine, Rohit Deshpandé, and Joshua D. Margolis, HBR, September 2011
- **"The Incentive Bubble,"** Mihir Desai, HBR, March 2012
- **"Managing Investors: An Interview with Sam Palmisano,"** Justin Fox, HBR, June 2014
- **"What Good Are Shareholders?"** Justin Fox and Jay W. Lorsch, HBR, July–August 2012

Originally published in May–June 2017. Reprint R1703B

The CEO View: Defending a Good Company from Bad Investors

A conversation with former Allergan CEO David Pyott by Sarah Cliffe

David Pyott had been the CEO of Allergan for nearly 17 years in April 2014, when Valeant Pharmaceuticals and Pershing Square Capital Management initiated the hostile takeover bid described in the accompanying article "The Error at the Heart of Corporate Leadership." He was the company's sole representative during the takeover discussions. When it became clear that the bid could not be fended off indefinitely, Pyott, with his board's blessing, negotiated a deal whereby Allergan would be acquired by Actavis (a company whose business model, like Allergan's, was growth oriented).

HBR: ***Would you describe Allergan's trajectory in the years leading up to the takeover bid?***

Pyott: We'd experienced huge growth since 1998, when I joined as just the third CEO of Allergan and the first outsider in that role. We restructured when I came in and again 10 years later, during the recession. Those cuts gave us some firepower for investing back into the economic recovery. After the recession we were

telling the market to expect double-digit growth in sales revenue and around the mid-teens in earnings per share.

Your investor relations must have been excellent.

They were. I am extremely proud to say that we literally never missed our numbers, not once in 17 years. We also won lots of awards from investor-relations magazines. You don't run a business with that in mind, but it's nice to be recognized.

In their article, Joseph Bower and Lynn Paine describe how difficult it is for any company to manage the pressure from investors who want higher short-term returns. You seem to have managed that well—until Valeant showed up. How?

Both buy-side and sell-side investors are like any other customer group. You should listen to what they say and respond when you can. But remember: Asking is free. If they say, "Hey, we want more," you have to be willing to come back with "This is what we can commit to. If there are better places to invest your funds, then do what you need to." Fortunately or unfortunately, I'm very stubborn.

Permit me a naive question: Since Allergan was going strong, why did it make sense to Valeant/Pershing Square to take you over and strip you down? I get that they'd make a lot of money, but wouldn't fostering continued growth make more in the long run?

Different business models. Valeant was a roll-up company; it wasn't interested in organic growth. Michael Pearson [Valeant's CEO] liked our assets—and he needed to keep feeding the beast. If he didn't keep on buying the next target, then the fact that he

was stripping all the assets out of companies he'd already bought would have become painfully obvious.

He couldn't do it alone, given his already weak balance sheet, so he brought Ackman in—and Pershing Square acquired 9.7% of our stock without our knowledge. This was meant to act as a catalyst to create a "wolf pack." Once the hedge funds and arbitrageurs get too big a position, you lose control of your company.

I still thought we had a strong story to tell—and I hoped I could get long-term-oriented shareholders to buy new stock and water down the hedge funds' holdings. But almost nobody was willing to up their position. They all had different reasons—some perfectly good ones. It was a lesson to me.

That must have been disappointing.

Yes. It's poignant—some of those same people say to me now, "We miss the old Allergan. We're looking for high-growth, high-innovation stocks and not finding them." I just say, "I heartily agree with you."

Another thing that surprised and disappointed me was that I couldn't get people who supported what we were doing— who understood why we were not accepting the bid, which grossly undervalued the company—to talk to the press. Several people said they would, but then folks at the top of their companies said no. And the reporters who cover M&A don't know the companies well. The people who cover pharma are deeply knowledgeable—but once a company is in play, those guys are off the story day-

to-day. So the coverage was more one-sided than we'd have hoped for.

Is the trend toward activist investors something that the market will eventually sort out?

Activist and hostile campaigns have been propelled by extraordinarily low interest rates and banks' willingness to accept very high leverage ratios. Recently investor focus has returned to good old-fashioned operational execution by management. But I do think that investment styles go in and out of fashion. I never would have guessed that when I went to business school.

Do you agree with Bower and Paine that boards and CEOs need to focus less on shareholder wealth and more on the well-being of the company?

Look at it from a societal point of view: A lot of the unrest we've seen over the past year is rooted in the idea that wealthy, powerful people are disproportionately benefiting from the changes happening in society. A lot of companies think that they need to make themselves look more friendly, not just to stockholders but to employees and to society. Having a broader purpose—something beyond simply making money—is how you do that and how you create strong corporate cultures.

I don't believe that strong performance and purpose are at odds, not at all. My own experience tells me that in order for a company to be a really high performer, it needs to have a purpose. Money matters to employees up to a point, but they

want to believe they're working on something that improves people's lives. I've also found that employees respond really favorably when management commits to responsible social behavior. I used to joke with employees about saving water and energy and about recycling: "Look, I'm Scottish, OK? I don't like waste, and it saves the company money." That's a positive for employees.

Did that sense of purpose pay off when you were going through the takeover bid?

Absolutely. I left day-to-day operations to our president, Doug Ingram, that year. And we grew the top line 17%—more than $1 billion—the best operating year in our 62-year history. I remember an R&D team leader who came up to me in the parking lot and said, "Are you OK? Is there anything I can do?" I answered him, "Just do your job better than ever, and don't be distracted by the rubbish you read in the media." Employees all over the world outdid themselves, because they believed in the company.

What changes in government rules and regulations would improve outcomes for the full range of stakeholders?

My favorite fix is changing the tax rates. Thirty-five percent is woefully high relative to the rest of the world. If we got it down to 20%, we'd be amazed at how much investment and job creation happened in this country. The high rates mean that we're vulnerable to takeovers that have tax inversion as a motivator. We were paying 26%, and Valeant [headquartered in

Canada] paid 3%. I think the capital gains taxes could be changed—in a revenue-neutral way—to incentivize holding on to stocks longer.

Shifting gears again: If a company wants to reorient itself toward long-term growth, what has to happen?

I think it's hard for a CEO to change his or her spots. Some can, but most can't. So in most cases you're going to need a new leader. And the board of directors really has to buy into it, because not only are you changing your strategy, you're changing your numbers. You must have a story to tell, for example: "For the next three years, we're not going to deliver 10% EPS growth. It's going to be 5% while we invest in the future. And that's not going to pay off until after three years, so you'll have to be patient." You have to be very, very clear about it.

And then everyone—the board, the investors, the lab technicians, the salespeople—will watch you to see if you're serious. It will take a lot of fortitude and determination. It's not impossible, but it's extremely difficult.

Originally published in May–June 2017. Reprint R1703B

Finally, Evidence That Managing for the Long Term Pays Off

by Dominic Barton, James Manyika, and Sarah Keohane Williamson

Companies deliver superior results when executives manage for long-term value creation and resist pressure from analysts and investors to focus excessively on meeting Wall Street's quarterly earnings expectations. This has long seemed intuitively true to us. We've seen companies such as Unilever, AT&T, and Amazon succeed by sticking resolutely to a long-term view. And yet we have not had the comprehensive data needed to quantify the payoff from managing for the long term—until now.

New research, led by a team from McKinsey Global Institute in cooperation with FCLT Global, found that companies that operate with a true long-term mindset have consistently outperformed their industry peers since 2001 across almost every financial measure that matters.

Firms focused on the long term exhibit stronger fundamentals and performance

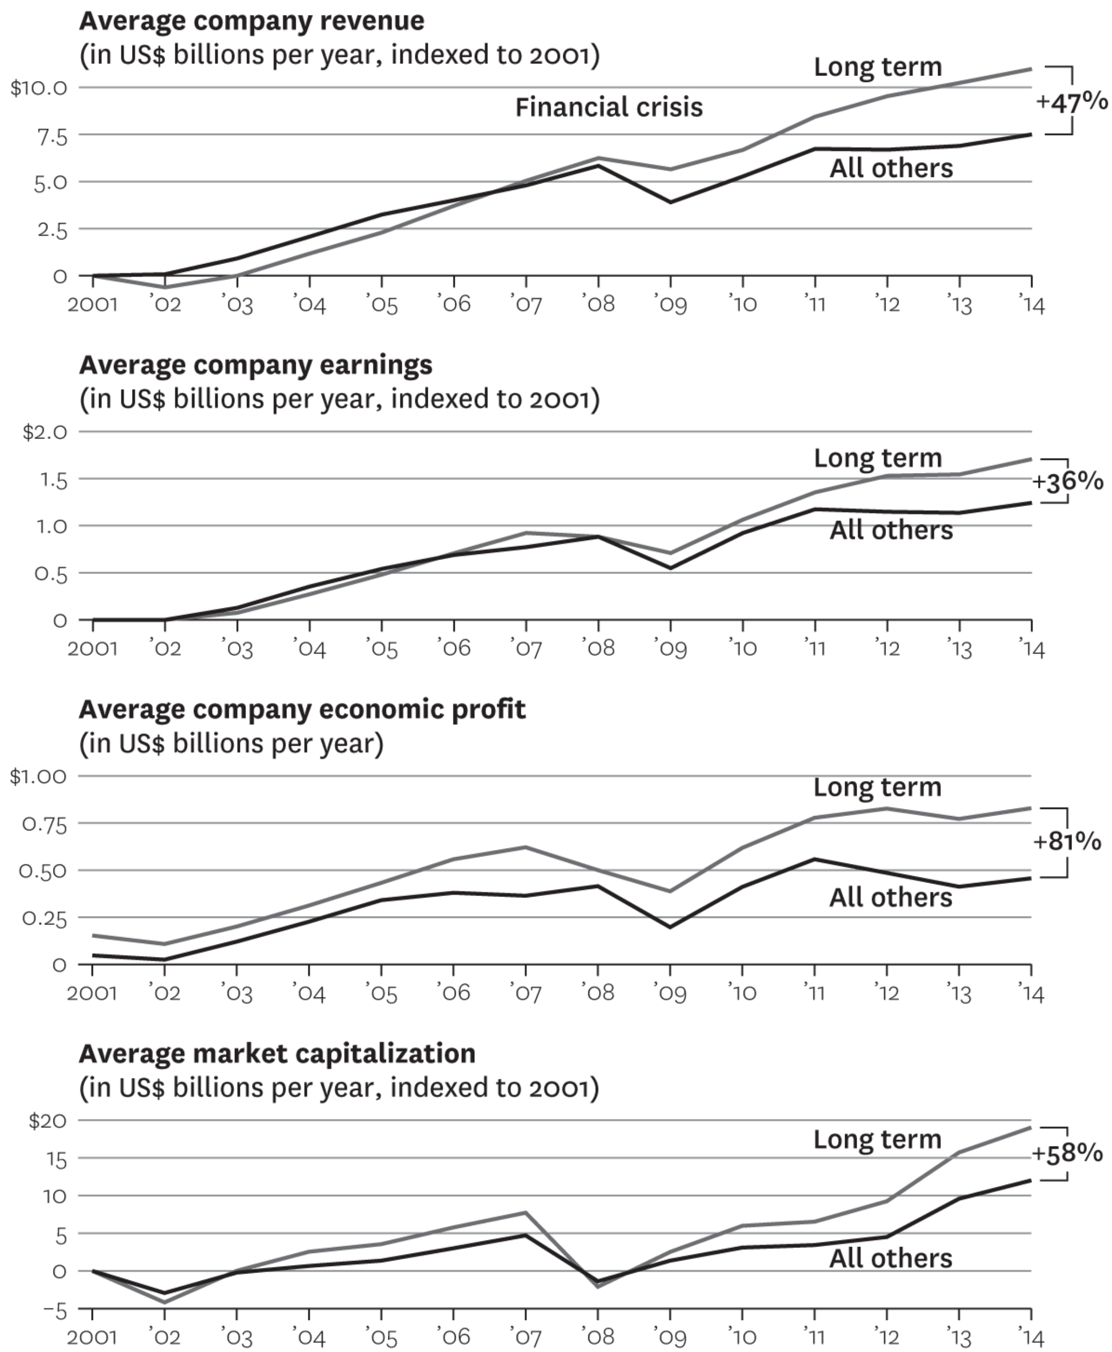

Source: McKinsey Global Institute.

The differences were dramatic. Among the firms we identified as focused on the long term, average revenue and earnings growth were 47% and 36% higher, respectively, by 2014, and

market capitalization grew faster as well. The returns to society and the overall economy were equally impressive. By our measures, companies that were managed for the long term added nearly 12,000 more jobs on average than their peers from 2001 to 2015. We calculate that U.S. GDP over the past decade might well have grown by an additional $1 trillion if the whole economy had performed at the level our long-term stalwarts delivered—and generated more than five million additional jobs over this period.

Who are these overachievers and how did we identify them? We'll dive into those answers shortly. But first, it's worth pausing to consider why finding conclusive data that establishes the rewards from long-term management has been so hard—and just how tangled the debate over this issue has been as a result.

In recent years we have learned a lot about the causes of short-termism and its intensifying power. We know from FCLT surveys, for example, that 61% of executives and directors say that they would cut discretionary spending to avoid risking an earnings miss, and a further 47% would delay starting a new project in such a situation, even if doing so led to a potential sacrifice in value. We also know that most executives feel the balance between short-term accountability and long-term success has fallen out of whack; 65% say the short-term pressure they face has increased in the past five years. We can all see what appear to be the results of excessive short-termism in the form of record

levels of stock buybacks in the U.S. and historic lows in new capital investment.

But while measuring the increase in short-term pressures and identifying perverse incentives is fairly straightforward, assessing the ultimate impact of corporate short-termism on company performance and macroeconomic growth is highly complex. After all, “short-termism” does not correspond to any single quantifiable metric. It is a confluence of so many complex factors it can be nearly impossible to pin down. As a result, despite persistent calls for more long-term behavior from us and from CEOs who share our views, such as Larry Fink of BlackRock and Mark Wiseman, the former head of the Canada Pension Plan Investment Board, a genuine debate has continued to rage among economists and analysts over whether short-termism really destroys value.

Academic studies have linked the possible effects of short-termism to lower investment rates among publicly traded firms and decreased returns over a multiyear time horizon. Ambitious work has even attempted to quantify economic growth foregone due to cuts in R&D expenditure driven by short-termism, putting it in the range of about 0.1% per year. Other researchers, however, remain skeptical. How, they ask, could corporate profits in the U.S. remain so high for so long if short-termism were such a drag on performance? And isn’t the focus on quarterly results a natural outgrowth of the rigorous corporate governance that keeps executives accountable?

What We Actually Measured—and the Limits of Our Knowledge

To help provide a better factual base for this debate, MGI, working with McKinsey colleagues from our Strategy & Corporate Finance practice as well as the team at FCLT Global, began last fall to devise a way to systemically measure short-termism and long-termism at the company level. It started with developing a proprietary Corporate Horizon Index. The data for this index was drawn from 615 nonfinance companies that had reported continuous results from 2001 to 2015 and whose market capitalization in that period had exceeded $5 billion in at least one year. (We wanted to focus on companies large enough to feel the potential short-term pressures exerted by shareholders, boards, activists, and others.) Collectively, our sample accounts for about 60%–65% of total U.S. public market capitalization over this period. To further ensure valid results and to avoid bias in our sample, we evaluated all companies in our index only relative to their industry peers with similar opportunity sets and market conditions and tracked them over several years. We also looked at the proportional composition of the long-term and short-term groups to ensure they are approximately equivalent, so that the differential performance of individual industries cannot bias the overall results, and conducted other tests and controls to ensure statistical robustness.

One final caveat: While we firmly believe our index enables us to classify companies as "long-term" in an unbiased manner, our

findings are descriptive only. We aren't saying that a long-term orientation causes better performance, nor have we controlled for every factor that could impact the relationship between those two. All we can say is that companies with a long-term orientation tend to perform better than similar but short-term-focused firms. Even so, the correlation we uncovered between behaviors that typify a longer-term approach and superior historical performance deliver a message that's hard to ignore.

To construct our Corporate Horizon Index, we identified five financial indicators, selected because they matched up with five hypotheses we had developed about the ways in which long- and short-term companies might differ. These indicators and hypotheses were:

- **Investment.** The ratio of capex to depreciation. We assume long-term companies will invest more and more-consistently than other companies.
- **Earnings quality.** Accruals as a share of revenue. Our belief is that the earnings of long-term companies will rely less on accounting decisions and more on underlying cash flow than other companies.
- **Margin growth.** Difference between earnings growth and revenue growth. We assume that long-term companies are less likely to grow their margins unsustainably in order to hit near-term targets.

- **Earnings growth.** Difference between earnings-per-share (EPS) growth and true earnings growth. We hypothesize that long-term companies will focus less on things like Wall Street's obsession with earnings-per-share, which can be influenced by actions such as share repurchases, and more on the absolute rise or fall of reported earnings.
- **Quarterly targeting.** Incidence of beating or missing EPS targets by less than two cents. We assume long-term companies are more likely to miss earnings targets by small amounts (when they easily could have taken action to hit them) and less likely to hit earnings targets by small amounts (where doing so would divert resources from other business needs).

After running the numbers on these indicators, two broad groups emerged among those 615 large and midcap U.S. publicly listed companies: a "long-term" group of 164 companies (about 27% of the sample), which were either long-term relative to their industry peers over the entire sample or clearly became more long-term between the first half of the sample period and the second half, and a baseline group of the 451 remaining companies (about 73% of the sample). The performance gap that subsequently opened between these two groups of companies offers the most compelling evidence to date of the relative cost of short-termism—and the real payoff that arises from managing for the long term.

Trillions of Dollars of Value Creation at Stake

To recap, from 2001 to 2014, the long-term companies identified by our Corporate Horizons Index increased their revenue by 47% more than others in their industry groups and their earnings by 36% more, on average. Their revenue growth was less volatile over this period, with a standard deviation of growth of 5.6%, versus 7.6% for all other companies. Our long-term firms also appeared more willing to maintain their strategies during times of economic stress. During the 2008–2009 global financial crisis, they not only saw smaller declines in revenue and earnings but also continued to increase investments in research and development while others cut back. From 2007 to 2014, their R&D spending grew at an annualized rate of 8.5%, greater than the 3.7% rate for other companies.

Another way to measure the value creation of long-term companies is to look through the lens of what is known as "economic profit." Economic profit represents a company's profit after subtracting a charge for the capital that the firm has invested (working capital, fixed assets, goodwill). The capital charge equals the amount of invested capital times the opportunity cost of capital—that is, the return that shareholders expect to earn from investing in companies with similar risk. Consider, for example, Company A, which earns $100 of after-tax operating profit, has an 8% cost of capital and $800 of invested capital. In this case its capital charge is $800 times 8%, or $64. Subtracting the capital charge from profits gives $36 of economic

profit. A company is creating value when its economic profit is positive, and destroying value if its economic profit is negative.

With this metric, the gap between long-term companies and the rest is even bigger. From 2001 to 2014 those managing for the long term cumulatively increased their economic profit by 63% more than the other companies. By 2014 their annual economic profit was 81% larger than their peers, a tribute to superior capital allocation that led to fundamental value creation.

No path goes straight up, of course, and the long-term companies in our sample still faced plenty of character-testing times. During the last financial crisis, for example, they saw their share prices take greater hits than their short-term counterparts. Afterward, however, the long-term firms significantly outperformed, adding an average of $7 billion more to their companies' market capitalization from 2009 and 2014 than their short-term peers did.

While we can't directly measure the cost of short-termism, our analysis gives an indication of just how large the value of what's being left on the table might be. As noted earlier, if all public U.S. companies had created jobs at the scale of the long-term-focused organizations in our sample, the country would have generated at least five million more jobs from 2001 and 2015—and an additional $1 trillion in GDP growth (equivalent to an average of 0.8 percentage points of GDP growth per year). Projecting forward, if nothing changes to close the gap between the long-term group and the others, then the U.S. economy could be

giving up another $3 trillion in foregone GDP and job growth by 2025. Clearly, addressing persistent short-termism should be an urgent issue not just for investors and boards but also for policy makers.

Where Do We Go from Here?

Our research is just a first step toward understanding the scope and magnitude of corporate short-termism. For instance, our initial dataset was limited to the U.S., but we know the problem is a global one. How do the costs and drivers differ by regions? Our sample set consists only of publicly listed companies. How do the effects we discovered differ among private companies or among public companies with varying types of ownership structures? Are there metrics that can help predict when a company is becoming too short-term—and how do they differ among industries? Most important, what are the interventions that will prove most effective in shifting organizations onto a more productive long-term path?

On this last point, we and many others have identified steps that executives, boards, and institutional investors can take to achieve a better balance between hitting targets in the short term and operating with a persistent long-term vision and strategy. These range from creating investment mandates that reward long-term value creation, to techniques for "de-biasing" corporate capital allocation, to rethinking traditional approaches to investor relations and board composition. We will return to

HBR in coming months with more data and insights into how companies can strengthen their long-term muscles.

The key message from this research is not only that the rewards from managing for the long term are enormous; it's also that, despite strong countervailing pressures, real change *is* possible. The proof lies in a small but significant subset of our long-term outperformers—14%, to be precise—that didn't start out in that category. Initially, these companies scored on the short-term end of our index. But over the course of the 15-year period we measured, leaders at the companies in this cohort managed to shift their corporations' behavior sufficiently to move into the long-term category. What were the practical actions these companies took? Exploring that question will be a major focus for our research in the coming year. For now, the simple fact of their success is an inspiration.

Further Reading

- **"Capitalism for the Long Term,"** Dominic Barton, *Harvard Business Review*, March 2011
- **"Focusing Capital on the Long Term,"** Dominic Barton and Mark Wiseman, *Harvard Business Review*, January–February 2014
- **"Where Boards Fall Short,"** Dominic Barton and Mark Wiseman, *Harvard Business Review*, January–February 2015

- **"The Short Long,"** Andrew G. Haldane and Richard Davies, speech, Bank of England, May 2011
- **"Profits Without Prosperity,"** William Lazonick, *Harvard Business Review*, September 2014
- **"Does a Long-Term Orientation Create Value?"** Caroline Flammer and Pratima Bansal, *Strategic Management Journal*, February 2017
- **"Businesses Can and Will Adapt to the Age of Populism,"** the *Economist*, January 2017

Originally published in February 2017. Reprint H03GCC

Now What?

by Joan C. Williams and Suzanne Lebsock

FAREWELL TO THE WORLD where men can treat the workplace like a frat house or a pornography shoot. Since Hollywood producer Harvey Weinstein was accused of sexual misconduct in early October, similar allegations have been made about nearly 100 other powerful people. They all are names you probably recognize, in fields including media, technology, hospitality, politics, and entertainment. It's a watershed moment for workplace equality and safety; 87% of Americans now favor zero tolerance of sexual harassment.

Not only is this better for women, but it's better for most men. A workplace culture in which sexual harassment is rampant is often one that also shames men who refuse to participate. These men-who-don't-fit, like the mistreated women, face choices about whether and how to intervene without endangering their careers.

Still, it's unnerving for many men to see the numbers of those toppled by accusations grow ever higher. The recent summary dismissals of high-powered executives and celebrities have

triggered worries that any man might be accused and ruined. Half of men (49%) say the recent furor has made them think again about their own behavior around women. Men wonder whether yesterday's sophomoric idiocy is today's career wrecker.

This is not a fight between men and women, however. One of the journalists to break the Weinstein story was Ronan Farrow, son of Mia Farrow and Woody Allen. Yes, that Woody Allen—the one who married his longtime girlfriend's daughter and is alleged to have sexually abused another daughter. "Sexual assault was an issue that had touched my family," said Farrow, who noted that this experience was instrumental in driving his reporting.

To repeat: This is not a fight between men and women. It's a fight over whether a small subgroup of predatory men should be allowed to interfere with people's ability to show up and do what they signed up for: work.

Several changes in the past decade have brought us to this startling moment. Some were technological: The internet enables women to go public with accusations, bypassing the gatekeepers who traditionally buried their stories. Other changes were cultural: A centuries-old stereotype—the Vengeful Lying Slut—was drained of its power by feminists who coined the term "slut shaming" and reverse-shamed those who did it. Just as important, women have made enough inroads into positions of power in the press, corporations, Congress, and Hollywood that they no longer have to play along with the boys' club; instead

they can, say, lead the charge to force Al Franken's resignation or break the story on Harvey Weinstein.

The result of all these changes is what social scientists call a norms cascade: a series of long-term trends that produce a sudden shift in social mores. There's no going back. The work environment now is much different from what it was a year ago. To put things plainly, if you sexually harass or assault a colleague, employee, boss, or business contact today, your job will be at risk.

How the #MeToo Movement Changes Work

As commonplace as these dismissals have come to seem, we know that we are only beginning to scratch the surface of the harassment culture. In "You Can't Change What You Can't See: Interrupting Racial & Gender Bias in the Legal Profession," a forthcoming study of lawyers conducted by the Center for WorkLife Law (which Joan directs) for the American Bar Association, researchers found sexual harassment to be pervasive. Eighty-two percent of women and 74% of men reported hearing sexist comments at work. Twenty-eight percent of women and 8% of men reported unwanted sexual or romantic attention or touching at work. Seven percent of women and less than 1% of men reported being bribed or threatened with workplace consequences if they did not engage in sexual behavior. Fourteen percent of women and 5% of men said that they had lost work opportunities because of sexual harassment,

which was also associated with delays in promotions, reduced access to high-profile assignments and sponsorship, bias against parents, and higher intent to leave. The three most acute types of harassment (excluding sexist remarks) were associated with reductions in income, demotions, loss of clients and office space, and removal from important committees.

Idea in Brief

When Hollywood producer Harvey Weinstein was accused of sexual harassment, the dam broke. Allegations of sexual misconduct were raised against many powerful people, and millions of women shared their own stories of harassment. It's a watershed moment for equality, say Williams, a legal scholar, and Lebsock, a feminist historian. Now 87% of Americans favor zero tolerance of harassment. Half of men are rethinking their own behavior. Over 75% of people are more likely to report sexist treatment at work. Everything has changed, for a simple reason: Women are being believed. Such was not the case in 1991, when Anita Hill claimed harassment by Supreme Court justice nominee Clarence Thomas. Back then women who came forward were often discredited as "vengeful, lying sluts." But that stereotype has been drained of power by feminists who coined the term "slut-shaming" and reverse-shamed those who did it. As the #MeToo and Time's Up movements demonstrate, women will no longer be silenced.

Translating outrage into action requires new norms of workplace conduct, which the authors outline. Firms are moving away from quiet settlements with victims and toward firing abusers. But employers still must follow due process and evaluate the credibility of reports. They need clear policies and fair procedures for handling harassment. No one's asking men to stop being men. But the reasonable assumption is that work relationships should be about work. You must not take one in a romantic direction if it's unwelcome, and the only way to safely tell what someone else wants is to ask. At the same time men shouldn't avoid women at work. That's unnecessary, unfair, and illegal: It deprives

women of opportunities simply because of their gender. Women, if colleagues make you uncomfortable, tell them. If you're harassed, report it. The authors aren't sure they'd have said that before #MeToo, but they do now, and it signals that the world has changed.

These patterns hold true beyond the legal profession. According to a recent study by researchers at Oklahoma State University, the University of Minnesota, and the University of Maine, women who were sexually harassed were 6.5 times as likely to change jobs as women who weren't. "I quit, and I didn't have a job. That's it. I'm outta here. I'll eat rice and live in the dark if I have to," remarked one woman in the study.

Low-wage women, who often live paycheck to paycheck, and women who are working in the U.S. illegally are the most vulnerable. A survey of nearly 500 Chicago hotel housekeepers revealed that 49% had encountered a guest who had exposed himself. Janitors who work the graveyard shift and farmworkers have had trouble defending themselves against predatory supervisors. And restaurant workers experience it from three directions. A 2014 report aptly titled "The Glass Floor," which shares the findings of a survey of 688 restaurant workers from 39 states, reveals that nearly 80% of the female workers had been harassed by colleagues. Nearly 80% had been harassed by customers, and 67% had been harassed by managers—52% of them on a weekly basis. Workers found customer harassment especially vexing because they were loath to lose crucial income from tips. Small wonder that almost 37% of sexual harassment

complaints filed by women with the Equal Employment Opportunity Commission in 2011 came from the restaurant industry.

The stories finally becoming public further highlight how sexual harassment subverts women's careers: Ashley Judd and Mira Sorvino found acting jobs harder to get after they rebuffed the voracious Weinstein. After Gretchen Carlson complained of a hostile work environment, she was assigned fewer hard-hitting interviews on *Fox & Friends* and, according to her legal complaint, was cut from her weekly appearances on the highly rated "Culture Warrior" segment of *The O'Reilly Factor*. Because word got out that Ninth Circuit judge Alex Kozinski sexually harassed clerks, many women did not apply for a clerkship at that court, which positions young lawyers to get clerkships at the U.S. Supreme Court—the biggest plum in the legal basket. When the ambitious congressional staffer Lauren Greene complained of sexual harassment by her boss, Representative Blake Farenthold, her career in politics evaporated. Today she works as a part-time assistant to a home builder.

A point often overlooked is that some sexual harassment victims are men. Men filed nearly 17% of sexual harassment complaints with the EEOC in 2016. Some men are harassed by women, but many are harassed by other men, some straight, some gay. A roustabout on an oil platform was harassed by coworkers on his eight-man crew, the U.S. Supreme Court found in 1998; the coworkers were offended by what they perceived as

his insufficient machismo. Recently the Metropolitan Opera suspended longtime conductor James Levine after several men accused him of masturbation-heavy abuse that took place from the late 1960s to the 1980s, when his victims were 16 to 20 years old.

Such behavior is no longer seen as a “tsking” matter. Historically, it has been hard to win a sexual harassment suit, but rapidly shifting public perceptions may change that. Seventy-eight percent of women say they are more likely to speak out now if they are treated unfairly because of their gender. About the same percentage of men (77%) say they are now more likely to speak out if they see a woman being treated unfairly. It’s a new day for a simple reason: Women are being believed.

Everything Is Changing

The strongest indicator that we’re experiencing a norms cascade came when Senate Majority Leader Mitch McConnell stood up for the women—four of them at the time—who had come forward with revelations about senatorial candidate Roy Moore.

“I believe the women,” McConnell said.

The statement stands in stark contrast to Anita Hill’s treatment in 1991, when she testified before the Senate Judiciary Committee that Clarence Thomas, then a nominee to the Supreme Court, had sexually harassed her. Senators subjected her to a humiliating inquisition, watched by a rapt national television audience. Another former employee was waiting in the

wings to describe how Thomas had sexually harassed her, too. But she was never called to testify. Instead, Hill withstood the all-male committee's bullying alone. After the hearings, opposition to Hill made her life at the University of Oklahoma so difficult that she left her tenured position—an object lesson on the risks facing anyone who dared to raise a charge of sexual harassment.

A recent poll by NPR dramatizes the sudden shift: 66% of Americans think that women who reported sexual harassment were generally ignored five years ago. Only 26% think that women are ignored today. When did we begin believing the women? What changed? And what are the implications for men?

We can trace the disbelief of—or at best, disregard for—women to the old stereotype we mentioned earlier, the one that holds women to be fundamentally irrational, vengeful, deceitful, and rampantly sexual.

An ancient version of this stereotype appears in Genesis, in which Eve commits the first sin and then drags Adam and the rest of humanity down with her for all time. Through the ages in Judeo-Christian tradition, authors expounded upon feminine evil. Among the most vivid prose stylists were two German friars, who in 1486 produced the classic book of witch lore *The Malleus Maleficarum* (or *The Hammer of Witches*). "What else is woman but a foe to friendship, an unescapable punishment, a necessary evil, a natural temptation, a desirable calamity, a domestic danger, a delectable detriment, an evil of nature, painted with

fair colours!" they wrote. More to the point for us, perhaps, is their claim that a woman "is a liar by nature."

Although by the 19th century more-positive images of women arose, the stereotype of the Vengeful Lying Slut was too useful to die. It was imposed on entire classes of women, notably African-American women, as scholars have amply documented, and on working-class women pressured into sex by bosses. It was used to ostracize and humiliate high schoolers who found themselves suddenly disparaged as "easy." Whenever men, and sometimes boys, exploited women—or often girls—the stereotype of the Vengeful Lying Slut supplied the words to justify their behavior: She wanted it/asked for it/had it coming.

The stereotype alas persists. It underlies men's fears that they, too, will be brought down by false allegations. Some men have become so frightened that they now refuse to meet (or to eat with) a female colleague alone. When Roy Moore was accused of sexual assault, his campaign said he was the victim of a "witch hunt." That response is a telling and time-honored way of discrediting victims.

The #MeToo and Time's Up movements show that women can no longer be silenced by threats of slut shaming. When a manager at Google told one of the female engineers who worked there, "It's taking all my self-control not to grab your ass right now," she tweeted it out to the world. In the first 24 hours after actress Alyssa Milano suggested that victims of harassment reply "me too" to a tweet in October, 12 million women made #MeToo

posts on Facebook. Instead of distancing themselves from those challenging sexual harassment, as might have happened in the past, actors and actresses wore black to the 2018 Golden Globes to signal their solidarity.

Translating outrage into action, however, requires moving beyond hashtags toward new norms of workplace conduct. It’s a precarious moment, and a lot could go wrong. Just think what might have happened if the *Washington Post*, with admirable rigor, had not uncovered the truth when a woman approached it with a dramatic but false accusation against Roy Moore. Her purpose? To snooker the Post into publishing a bogus story and to thereby cast doubt on all mainstream media reporting the claims against Moore. But so far so good, with early signs that workplaces are indeed changing.

Firing Is the New Settlement

In the past companies often quietly paid to settle sexual harassment complaints against high-powered miscreants and tried to limit the damage through nondisclosure agreements. Incidents at Fox gave rise to at least seven settlements (some against Fox, some against individuals at Fox). Weinstein reportedly paid out eight. Despite getting large payouts, the plaintiffs were the ones who were forced to leave their companies, and many suffered career interruptions.

Quiet settlements are now becoming harder to justify. The unceremonious firings and forced resignations of famous men

demonstrate that companies are moving away from that strategy. Settlements will likely continue in some circumstances, such as a first offense involving mild or ambiguous behavior or a situation that is consensual but violates company standards. But long strings of settlements in egregious cases will increasingly be seen as a breach of the directors' duty to the company. Boards of directors have never tolerated financial fraud and violations of the Foreign Corrupt Practices Act, and they are likely to adopt the same standards for harassment—firing without severance pay.

It's important to recognize that most of the firings have occurred at companies with sophisticated legal and HR departments, on the advice of counsel and with the involvement of senior management or the board or both. We should not assume that they are disclosing all the evidence they have. Companies have a strong motive not to release such evidence, lest the former employee use it as ammunition in a defamation or wrongful discharge suit. That's what companies do when they sack someone for cause, and that's what they are doing here.

Some worry that people will be fired too quickly and without due process. One point that's often overlooked: Due process isn't required of private employers, only public ones. What people are trying to insist on, quite properly, are fair procedures that uncover the truth. Companies should follow the same procedures they use when an employee has been accused of any type of serious misconduct. Typically, the employee is placed on

leave while an investigation is performed. In most cases, although not all, that's what has been happening with sexual harassment cases.

Credibility assessments are, of course, important. Women are human beings, and sometimes human beings—male and female—lie. That's why we need to apply the standard methods we always use to assess credibility. Those methods are flawed, but they are all we have; if they will do for every other context, they will do for sexual harassment, too.

As we enter this new era, here's a comforting thought from someone who has spent his life thinking about how to ferret out the truth, the prominent evidence scholar Roger Park (a colleague of Joan's). His observation about sexual harassment is this: "Men have a motive to do it and lie, whereas women don't have a motivation to lie, considering what an ordeal it is." Making even *true* allegations of sexual harassment has historically been a poor career move.

That provides some assurance that reports of harassment are truthful. So do large numbers of people with similar stories. At least 42 women have come forward with allegations against Weinstein, and at least 10 against Ken Friedman, the New York restaurateur. At least a dozen people have made accusations against Kevin Spacey. Those numbers lend credibility to the allegations.

Employers who want to set up processes for handling harassment can begin with the standard sexual harassment

policies. The Society for Human Resources has one; others are free online. Organizational training should spell out what's acceptable, which will vary from company to company. Some companies may want to add detail in light of recent events. Surprising as it sounds, some people seem to need a heads-up that porn, kissing, back rubs, and nudity are not appropriate at work.

How can this be? Here's a clue. At a dinner Judge Kozinski held with law clerks, he steered the conversation to the "voluptuous" breasts of a topless woman in a film, according to someone present. When one woman at the dinner reacted negatively, Kozinski responded that, well, he was a man.

Some men have an urgent need to preserve sexual harassment as a prerogative because, they feel, their manliness is at stake. But theirs is just one definition of manliness—a toxic and outdated one. It's time to move on.

The Workplace Today

Virtually all women and most men are now aligned against that toxic brand of masculinity. No one is asking men to stop being men or for people to stop being sexual beings. What's happened is that a small group of men are being required to abandon the stereotype that "real men" need to be unrelentingly sexual without regard to context or consent.

The not-unreasonable assumption is that work relationships should be about work. Some organizations have no-dating

policies for that reason. If yours doesn't, remember that you must not take a relationship with a colleague in a romantic or sexual direction if doing so is unwelcome. Whether you can ask a colleague out is the source of much anxiety, especially in all-consuming work environments where people date coworkers because they spend so much time on the job that there's little opportunity to meet anyone else.

The only way to safely tell what someone else wants is to ask that person. Some men seem to have trouble discerning whether a woman is interested; Charlie Rose and Glenn Thrush said that they thought their feelings were reciprocated when women who received their overtures say they were not. This is not an unsolvable problem. If she's a work colleague and you'd like her to be something more, here's what to do: Imagine telling a woman who's been your friend forever that you'd like to take the relationship in a different direction. Ask in a way that gives her a chance to say that she prefers to remain a friend. No harm, no foul. What if your work colleague says no when she really means yes? Well, then, she's got to live with that. Let her. Let her change her mind if she wants to.

We all know that deals and crucial networking happen over lunch, dinner, and drinks. Socializing in this manner is fine. But if you do socialize with work colleagues, you need to realize that you can't behave inappropriately. Roy Price resigned from his job as head of Amazon Studios after Isa Hackett, an Amazon producer, publicly accused him of repeatedly propositioning her

in a cab on the way to a work party, telling her, "You'll love my dick," and later at the gathering whispering "anal sex" loudly in her ear in the presence of others. Hollywood commentator Nellie Andreeva noted that in a post-Weinstein world Price's behavior would have hurt Amazon's ability to attract female showrunners and actors. He would have been "completely ostracized," an anonymous source told Andreeva.

Where the Law Draws the Line

by Joan C. Williams

Many of the sexual harassment cases making headlines involve criminal behavior. Sexual assault and related offenses are defined in different ways in different states. To take just one example, New York law prohibits "forcibly touching the sexual or other intimate parts of another person for the purpose of degrading or abusing such person or for the purpose of gratifying the actor's sexual desire." The statute helpfully adds that this "includes squeezing, grabbing or pinching." It should not be surprising or puzzling that such behavior is not acceptable.

Gentlemen, you already know not to invite a woman to discuss a job and then meet her wearing a bathrobe and expose yourself. Charlie Rose reportedly did that, and according to several women, Harvey Weinstein did that—and more. Masturbating in work contexts is not only unacceptable but illegal, yet that's what Louis C.K. did. One does not stick one's tongue down the throat of someone during a discussion of job prospects, as two women have claimed NPR's Michael Oreskes did. You do not kiss a colleague and lick her, as actor Andy Dick did.

Men do not expect to report to work and have their crotch grabbed. Women don't, either. It should not be frightening or confusing to be told this. But we understand why men are scared: Most sexual harassment

does not involve sexual assault, and if you've ever told an off-color joke at work, asked out a colleague, or maybe been a little handsy at a holiday party (or know someone who has), we bet you've been thinking a lot about sexual harassment lately.

Employment law has a sober and balanced approach that is fully protective of the rights of men accused of sexual harassment. It defines two kinds of sexual harassment:

Quid Pro Quo Harassment

Making sexual favors a condition of any workplace opportunity is illegal under federal law. To win a lawsuit alleging it, a woman has to prove that someone with authority over her threatened to take a negative employment action unless she engaged in a sexual behavior—or promised her a promotion, raise, or other benefit if she did. Congressman John Conyers paid thousands of dollars to a staffer who said she was fired for refusing his sexual advances. According to accounts published in *New York* magazine, Roger Ailes tied women's work prospects to sex again and again: "If you want to play with the big boys, you have to lay with the big boys," he told a woman seeking a contract with the Republican National Committee in 1989. "No girls get a job here unless they're cooperative," he is reported to have said to a frightened 19-year-old in the sixties after he grabbed her and forcibly kissed her. Fifty years later, TV host Gretchen Carlson says, Ailes demoted and ultimately fired her for refusing to have sex with him. Fox settled Carlson's harassment case for $20 million. But it extended Bill O'Reilly's contract after O'Reilly settled a sexual harassment case against him for $32 million. Cases do not settle for that kind of money unless something has gone seriously wrong.

For men worried about *quid pro quo harassment*, the simplest approach is not to date someone you supervise. If you do, make sure it's consensual and remember that whether you stay together or break up, with respect to workplace issues you need to behave *exactly* as you would have if you'd never dated her. If you can't do that, don't date people you supervise. All this applies, of course, not only to men but also to women.

A Hostile Work Environment

Here again the legal test is quite protective of those accused of sexual harassment. To meet the legal definition the conduct must be unwelcome, the environment must be one that a reasonable person would consider hostile, the plaintiff herself must feel it to be hostile, and the behavior that makes the environment hostile must be severe or pervasive.

Moreover, plaintiffs very rarely win hostile environment cases that are based on a single "severe" incident. Almost invariably, they need to prove the behavior was "pervasive." How pervasive?

Very. In a 1993 Supreme Court case, a woman's boss made such comments as "You're a woman, what do you know?," "We need a man as the rental manager," and "Dumb-ass woman," and asked her in front of coworkers if she wanted to "go to the Holiday Inn" to negotiate a raise. He asked women to retrieve coins from the front pockets of his pants and threw objects on ground and asked women to pick them up. When the plaintiff complained to the boss about his conduct, he said he was joking and promised to stop, but he didn't. She quit, sued, and won: She had made it clear the behavior was unwelcome and that it personally offended her. The court found that a reasonable person would have felt the environment was hostile and that the hostility was pervasive.

This is so far beyond what most men would ever imagine is appropriate that they have little to fear. Still, the requirement that the plaintiff prove that she herself felt an environment to be hostile adds another layer of protection. So women need to speak up to demonstrate that they're personally experiencing that feeling, not just to show that an advance is unwelcome. "That makes me uncomfortable. We are at work" is enough.

It's OK to socialize with and date colleagues. But the law regarding "retaliation" requires that a colleague must be able to decline an invitation without consequences. The easiest way for a company to lose a sexual harassment suit is for a plaintiff to prove that she turned down

an advance or complained of a hostile environment and then suffered retaliation.

A new textbook example comes from Uber. Susan Fowler's manager propositioned her on the company chat site. She took a screenshot of it and went to HR, which gave her a choice: Either she could transfer out of her team, or she could stay. But if she stayed, the people in HR said, her boss might give her a poor performance review and they could do nothing about it.

Here's the problem: Giving someone a poor review because she turned you down is retaliation, which is illegal. Fowler didn't want to transfer, because her team's project was a great match for her specific skills, but she decided to do so. Fortunately, she found other great work to do, but unfortunately, she continued to encounter sexist behavior, which she reported to HR. One day her boss called her in and threatened to fire her if she didn't stop making complaints. That's retaliation too—and it's illegal.

You can still compliment your colleagues. But there's a big difference between "I like that dress" and "You look hot in that dress." What if she really *does* look hot? Remember, she signed up to be your colleague, not your girlfriend. Treat her like a colleague unless by mutual consent, you change your relationship.

Don't let the pendulum swing too far the other way and bizarrely avoid women completely. That's unnecessary, unfair, and illegal: It deprives women of opportunities simply because they are women. You cannot refuse to have closed-door meetings with women unless you refuse to have closed-door meetings with men. Otherwise women will be denied access to

all the sensitive information that's shared only behind closed doors, and that's a violation of federal law.

Moving forward, male allies will continue to play an important role in fighting harassment: If you see something, say something. It does take courage, but you can use a light touch. If you're standing around with a bunch of guys and a female colleague walks by, only to have someone say, "Wow, she's hot," you can say simply: "I don't think of her that way. I think of her as a colleague, and that's the way I suspect she'd like to be thought of."

Clear takeaways emerge for women, too. If a coworker tries to take a work relationship in a sexual direction, tell him clearly if that's unwelcome. If you face sexual joking that's making you uncomfortable, say, "This is making me uncomfortable" and expect it to stop. If you want to shame or jolly someone out of misbehavior while preserving your business relationships, consult Joan's *What Works for Women at Work*. Here's an approach that worked for one woman whose colleague proposed an affair: "I know your wife. She's my friend. You're married. There is just no way I would ever consider that. So let's not go there again."

But it's our final piece of advice that signals the tectonic shift: If you are being sexually harassed, report it. We're not sure if we would have advised that, in such a blanket and unnuanced way, even a year ago.

What we're seeing today is not the end of sex, or of seduction, or of *la différence*. What we're seeing is the demise of a work culture where women must submit to being treated, insistently and incessantly, as sexual opportunities. Most people, when they go to work, want to work. And now they can.

Originally published in January 2018. Reprint BG1801

The Omissions That Make So Many Sexual Harassment Policies Ineffective

by Debbie S. Dougherty

Our research began with a simple question: If 98% of organizations in the United States have a sexual harassment policy, why does sexual harassment continue to be such a persistent and devastating problem in the American workplace? As evidenced by recent headlines regarding ongoing sexual harassment in the National Park Service, Uber, and Fox News, it seems clear that sexual harassment policies have not stopped the problem they were designed to address.

Two bodies of research provided us with a possible direction as we explored the relationship between sexual harassment policies

and outcomes. First, scholars convincingly argue that sexual harassment is embedded in organizational culture. In other words, sexual harassment serves an important cultural function for some organizations. And as any executive who has tried to lead cultural change knows, organizational culture can be immutable.

Second, organizational cultures are embedded in a larger national culture in which men have traditionally been granted privileges over women. It does not take a deep analysis to recognize this truth. Women are typically paid less, regardless of education, qualifications, or years of service. There are more CEOs named John leading big companies than there are female CEOs. The male-centric nature of our national culture is so pervasive that even many women are male-centered, aligning themselves with men and masculinity to tap into male privilege while attempting (usually unsuccessfully) to avoid the disadvantaged space that women occupy in the workplace.

All of this means that both men and women can react to sexual harassment by blaming other women for "making trouble" or "putting up with bad behavior," or by suggesting that the sexually harassed women should quit, without considering that perhaps the perpetrators instead of their targets should leave the organization. These attitudes have real consequences. Consider: In the Fox News harassment case, the alleged perpetrators received larger settlements than the targets. Cultures of sexual

harassment are thus legitimized by drawing on the larger cultural imperative that privileges men over women.

Into this fraught cultural morass enters a well-intentioned document: the sexual harassment policy. To see how employees interpreted these policies, my colleague Marlo Goldstein Hode and I gave 24 employees of a large government organization a copy of the organization's sexual harassment policy, asking them to read it and then tell us about the policy. We asked them to talk about the policy in groups, and then we interviewed them individually.

We found that the actual words of the sexual harassment policy bore little resemblance to the employees' interpretations of the policy. Although the policy clearly focused on *behaviors* of sexual harassment, the participants almost universally claimed that the policy focused on *perceptions* of behaviors. Moreover, although the policy itself made clear that harassing behaviors were harassment regardless of either the gender or sexual orientation of the perpetrator or target, the employees focused almost exclusively on male-female heterosexual harassment. This shift is subtle but significant. For the participants, the policy was perceived as threatening, because any behavior could be sexual harassment if an irrational (typically female) employee perceived it as such. In this somewhat paranoid scenario, a simple touch on the arm or a nonsexual comment on appearance ("I like your hairstyle") could subject "innocent" employees (usually heterosexual males) to persecution as stipulated by the

policy. As a result, the organization’s sexual harassment policy was perceived as both highly irrational and as targeting heterosexual male employees. The employees shifted the meaning of the policy such that female targets of sexual harassment were framed as the perpetrators and male perpetrators were framed as innocent victims.

To accomplish this shift in meaning, the employees drew on assumptions of women being irrational and highly emotional and on assumptions of men being rational and competent. Through this intertwining of organizational policy, organizational culture, and national culture, the employees inverted the meaning of the sexual harassment policy, making it an ineffective tool in the fight against predatory sexual behavior in the workplace.

How can organizations combat the reinterpretation of sexual harassment policies? This question takes on urgency when we recognize that sexual harassment policies are table stakes in successfully managing the damaging behavior.

Remember that sexual harassment policies are not just legal documents. They are also culturally important, *meaning-making* documents that should play a role in defining, preventing, and stopping sexual harassment in an organization. The findings from our study suggest very specific language that may be useful in sexual harassment policies:

- **Include culturally appropriate, emotion-laden language in sexual harassment policies.** Our findings suggest that if you don’t

add this language, organizational members will include their own. For example, adding language such as “Sexual harassment is a form of predatory sexual behavior in which a person targets other employees” frames the behavior such that alternative interpretations may be more difficult to make. Using terms such as “predatory” instead of “perpetrator” and “target” instead of “victim” can shape how organizational members interpret the policy. Although policies tend to be stripped of emotions, it is essential for policy creators to recognize that policy creation is one of the most emotion-laden activities that organizational leaders are asked to accomplish. Because sexual harassment is such an emotionally laden topic, the creation of sexual harassment policies becomes even more emotionally challenging.

- **Sexual harassment policies should include bystander interventions as a required response to predatory sexual behavior.** Most policies place responsibility for reporting harassment exclusively on the target, which puts them in a vulnerable position. If they report the behavior, then they are likely to be viewed with suspicion by their colleagues, often becoming socially isolated from their coworkers. On the other hand, if they do not report the sexual harassment, then it is likely to continue unabated, creating harm for the targeted employee, and wider organizational ills, too. Mandating bystander intervention can relieve the target of

their sole responsibility for reporting and stopping predatory sexual behavior, and rightly puts the responsibility of creating a healthier organizational culture on *all* members of the organization.

Sexual harassment is complicated. If it were a simple problem involving just two people, we would have resolved the issue decades ago. But sexual harassment is a complicated, entrenched problem. Systems theory tells us that solutions need to match the complexity of the problem. Writing a policy is complicated, as our study showed. But it's also just a start. No policy, no matter how well crafted, will prevent sexual harassment on its own, nor will it change a culture of sexual harassment. A policy is a first step that needs to be followed by persistent training, a willingness to listen to targets, and a readiness to fire employees who prey sexually on other employees—regardless of how important the predator may be in the organization.

Originally published in May 2017. Reprint H03ONZ

How Do Your Workers Feel About Harassment? Ask Them

by Andrea S. Kramer and Alton B. Harris

If your business is serious about eliminating the risk of sexual harassment—and it should be—you need to approach the problem comprehensively. This means recognizing that sexual harassment is part of a continuum of interconnected behaviors that range from gender bias to incivility to legally actionable assault. All these kinds of misconduct should be addressed collectively, because sexual harassment is far more likely in organizations that experience offenses on the “less severe” end of the spectrum than in those that don’t.

There’s no one-size-fits-all program for eliminating inappropriate gender-related behaviors; the best programs specifically address the characteristics of each workplace’s culture. The vital first step, then, is to get an accurate picture of yours. How? Ask your employees directly. Do they see disparities in career opportunities? Are colleagues or supervisors rude to each other? Is there inappropriate sexual conduct? Do employees feel uncomfortable or unsafe at work?

The best way to find all this out is with a carefully designed employee survey. In this article we’ll offer some key principles for fashioning one, along with a model survey that you can adapt (which incorporates some of the recommendations the EEOC made for surveys in its 2017 proposed enforcement guidance on harassment). Our advice is based on insights we developed while working with major business organizations and conducting several hundred gender-focused workshops and moderated conversations around the United States.

Though we think a workplace climate survey can be immensely valuable, we caution that managers and leaders should proceed only if they're fully committed to thoroughly and quickly addressing inappropriate behavior in their organizations. Once the surveys are undertaken, they'll create expectations of remedial action. They might also attract unwanted publicity or be used against the company in future litigation. Those risks, however, are substantially outweighed by the opportunity to develop a work environment that's free of sexual misconduct, gender bias, and incivility.

Step 1: Communicate with Employees

Inform your employees that you're undertaking an effort to understand how fair, courteous, and safe their workplace is. The goal is to encourage engaged and completely candid answers to the survey. For that reason, it should be anonymous and administered by a third party, not your HR department. Employees won't have faith that their answers are confidential if the survey is conducted in-house, and if you don't offer true anonymity, their responses are less likely to be honest.

It's crucial for employees to believe that management considers an unbiased and harassment-free workplace a priority and is sincere in its commitment to that objective. That will happen only if senior management openly endorses the initiative, communicates the importance of supporting it to the

entire management team, and periodically speaks to all employees about it.

Employees also need to believe that the end result will be better policies for everyone. This last point can't be emphasized too strongly. If the steps you take to combat inappropriate gender-related behaviors are seen solely as efforts to "protect" women because of their vulnerability, the initiative will backfire.

The first organization-wide letter to employees might begin with a statement like this:

> *We are gathering information on a confidential basis to better understand our business's workplace climate. Our goal is to ensure that all employees receive equal opportunities, respect, and resources in a workplace that is free of incivility and does not tolerate inappropriate sexual conduct.*
>
> *The survey that you'll receive shortly is the first step toward achieving that objective.*
>
> *We have hired a third-party administrator to conduct the survey on a strictly anonymous basis. Your answers and identity will be carefully protected from disclosure.*
>
> *The administrator will contact you separately and detail its procedures for preserving anonymity.*
>
> *The survey you'll receive is divided into four parts: gender bias, incivility, inappropriate sexual conduct, and overall workplace climate. All four areas are important, so please be as candid as possible in giving your views.*

Employees should also be told that only the third-party administrator will see the raw survey results and that it will provide an analysis of those results to management. Once management receives that report, employees should be advised of the nature of and timetable for next steps.

We suggest that you emphasize that because the survey is anonymous, your organization cannot investigate or remedy specific claims raised by respondents—unless the incidents are separately reported in accordance with existing company procedures. Urge your employees to use those procedures if appropriate.

Step 2: Draw Up Your Survey

Whether you start with the assessment that we provide in this article or create your own questions, you should tailor your survey to your organization's culture and climate. Keep in mind the following:

- Avoid questions that could be used—or thought to be used—to identify participants, such as those about title, age, tenure with the company, responsibilities, and workplace location.
- Don't ask about marital or domestic status, sexual preference, children, or prior involvement in sexual misconduct investigations or proceedings. An inappropriate

question in a job interview is equally inappropriate in a workplace climate survey.

- Keep the survey on point. Resist the temptation to use it as an opportunity to ask employees more broadly about their experiences, expectations, and future plans.
- Make the survey short and unambiguous. It should take no more than 10 minutes to finish. You may use true/false, multiple choice, or open-ended questions, but in our experience, the most useful approach is to incorporate a scale. Develop a series of statements that participants will be asked to indicate their degree of agreement with, using a scale of 1 (strongly disagree) to 7 (strongly agree). With statements that are intended to examine the frequency of specific behaviors, use a scale of 1 (very frequently) to 4 (never).

Step 3: Evaluate

A workplace climate survey needs no statistical evaluation beyond a simple tabulation. You're just attempting to determine whether some of your employees believe there are gender-related problems in your work environment and what those problems are.

Bear in mind that the survey is not an end in itself; it's a tool to identify whether you need new policies, practices, and procedures to eliminate inappropriate behavior and protect your

employees against sexual harassment. Your results may indicate additional steps are necessary. You might need to assemble focus groups, conduct personal interviews, or host roundtable discussions. Since your goal is to ensure you have a welcoming, supportive, and productive workplace, the real work begins once you have a clear picture of your business's actual climate. Here is a template you can use when constructing your survey:

Model Workplace Climate Survey

Complete the following survey about your experience at XYZ Company, without referring to experiences at any prior organizations. The value of this survey depends directly on getting an accurate view of our workplace culture, so please answer all questions as honestly as possible.

1. Which of the following describes your gender?

- Male
- Female
- Prefer to self-describe (specify)
- Prefer not to say

Gender Bias

2. I feel valued by the organization.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

3. I believe my opportunities for career success are negatively affected by my gender.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

4. The people I work with treat me with respect and appreciation.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

5. My views are encouraged and welcomed by my supervisors and senior leaders without regard to my gender.

   (1) Strongly disagree

   (2) Disagree

   (3) Slightly disagree

   (4) Neither agree nor disagree, or have no opinion

   (5) Slightly agree

   (6) Agree

   (7) Strongly agree

6. Career-enhancing assignments and opportunities are disproportionately given to men.

   (1) Strongly disagree

   (2) Disagree

   (3) Slightly disagree

   (4) Neither agree nor disagree, or have no opinion

   (5) Slightly agree

   (6) Agree

   (7) Strongly agree

Civility

7. My coworkers are courteous and friendly.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

8. I'm aware of unpleasant and negative gossip in the workplace.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

9. I'm aware of abusive, disrespectful, or hostile treatment of employees.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

10. I'm aware of bullying behavior in the workplace.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

11. There are adverse consequences for senior leaders who are abusive, disrespectful, or hostile.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

12. I have been criticized for my personal communication style or appearance.

(1) Very frequently

(2) Somewhat frequently

(3) Not at all frequently

(4) Never

13. All individuals are valued here.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

Inappropriate Sexual Conduct

14. I have experienced or witnessed unwanted physical conduct in the workplace or by coworkers away from the workplace.

(1) Very frequently

(2) Somewhat frequently

(3) Not at all frequently

(4) Never

15. I have witnessed or heard of offensive or inappropriate sexual jokes, innuendoes, banter, or comments in our workplace.

    (1) Very frequently

    (2) Somewhat frequently

    (3) Not at all frequently

    (4) Never

16. I have witnessed or heard of the electronic transmission of sexually explicit materials or comments by coworkers.

    (1) Very frequently

    (2) Somewhat frequently

    (3) Not at all frequently

    (4) Never

17. I have received sexually inappropriate phone calls, text messages, or social media attention from a coworker.

    (1) Very frequently

    (2) Somewhat frequently

    (3) Not at all frequently

(4) Never

18. I have been asked or have witnessed inappropriate questions of a sexual nature.

(1) Very frequently

(2) Somewhat frequently

(3) Not at all frequently

(4) Never

19. I have been the subject of conduct that I consider to be sexual harassment.

(1) Very frequently

(2) Somewhat frequently

(3) Not at all frequently

(4) Never

20. Managers here tolerate or turn a blind eye to inappropriate sexual conduct.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

21. I feel unsafe at work because of inappropriate sexual conduct by some individuals.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

22. I've seen career opportunities be favorably allocated on the basis of existing or expected sexual interactions.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

23. I would be comfortable reporting inappropriate sexual conduct by a coworker.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

24. I would be comfortable reporting inappropriate sexual conduct by a supervisor.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

Overall Workplace Climate

25. My productivity has been affected by inappropriate gender-related behavior in the workplace.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

26. I have considered leaving my job because of inappropriate gender-related behavior in the workplace.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

27. Star performers are held to the same standards as other employees with respect to inappropriate gender-related behavior.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

28. I have experienced or witnessed inappropriate gender-related behavior by third parties (such as customers, vendors, and suppliers) associated with our organization.

(1) Very frequently

(2) Somewhat frequently

(3) Not at all frequently

(4) Never

29. The organization's policies and processes with respect to prohibiting and reporting inappropriate gender-related behavior are easy to understand and follow.

(1) Strongly disagree

(2) Disagree

(3) Slightly disagree

(4) Neither agree nor disagree, or have no opinion

(5) Slightly agree

(6) Agree

(7) Strongly agree

Originally published in January 2018. Reprint BG1801

Getting Men to Speak Up

by Michael S. Kimmel

In early November 1991, a month after Anita Hill's testimony about being sexually harassed by Supreme Court nominee Clarence Thomas, my mother invited me to dinner. After a long and pleasant meal, she told me that Hill's stories were all too familiar. When my mother was in graduate school, her mentor groped her. She left school the next day and didn't complete her PhD for 30 years.

Back in the 1990s, Hill wasn't believed when she bravely came forward. Instead she was vilified by the Senate Judiciary Committee as a woman scorned, as "a little bit nutty and a little bit slutty," as a now-contrite David Brock put it in his article

smearing Hill. That response set the tone: Over the next 25 years, whenever a woman stood up to publicly accuse men like Bill Cosby or Bill Clinton of sexual assault, she usually ended up being the one on trial in the court of public opinion, charged with a lack of credibility.

But outside this public narrative, something started to shift: Women like my mother began to speak privately about their painful experiences. Mothers told their children, wives told their husbands, women told their friends, daughters told their parents. And they were believed.

Social scientists who study movements often speak of the three elements of revolution. First come the structural preconditions—long-term institutional changes that slowly build pressure, sometimes without even being noticed. In this case, those 25 years of simmering private conversations paved the way for today's widespread backlash against harassment. The second element of a revolution is precipitants—pivotal events that cause change to rapidly accelerate. One precipitant here was the 2016 release of the *Access Hollywood* videotape of Donald Trump bragging about kissing and groping women. After his election to the U.S. presidency despite this evidence, many women were both incredulous and furious.

Finally, there are trigger events that ignite a major explosion. In this case it was the rapid succession of revelations about Roger Ailes, Bill O'Reilly, and Harvey Weinstein. In what seemed like a first, the women's tales of abuse were not doubted—they were

believed. And so #MeToo began, a reckoning so public that the women who spoke out were named Time magazine's people of the year in 2017.

We are in a new moment. For many of us, particularly men, it is scary and uncomfortable. Men are feeling vulnerable and afraid of false accusations (or perhaps true ones). They fear that things they did a long time ago will be reevaluated under new rules. They tell me they're walking on eggshells. Because of this, many men are staying silent rather than taking part in the conversation. And yet inaction isn't necessarily the right approach; there are important things men can do and say to support the women in their lives.

My experience studying masculinity and working with companies on sexual harassment has led me to focus on how men can take action to address this problem in the workplace. To do so effectively, we must come to terms with four questions: Why do men harass women? Don't they know it's wrong? How do they get away with it? And finally, what can we do about it?

Why Do Men Harass Women?

This one's easy. Men do it because they feel they can. It's hardly the case that men are so overcome by lust that they cannot restrain themselves, as some people have suggested. No, it's often about being in a position of power and feeling entitled to have access to women. These male harassers are emboldened to

act by their privilege and authority and by the fact that their targets are in a weaker and more vulnerable position.

Don't They Know It's Wrong?

Nearly all of us know that grabbing a woman by her genitals, patting her butt, making lewd comments, or forcing her to engage in sexual activity is wrong. This is not some blurry line we have to negotiate. *We know*. "They let you do it" is the most telling quote from that *Access Hollywood* tape. Trump is saying, in effect, *You see what a big celebrity I am? Look what I can get away with*.

Some men, however, may not realize that the occasional shoulder massage, calling women "sweetie" or "honey," or making suggestive comments is also wrong. Men who are older tend to fall into this category. It's startling to remember that a mere two generations ago, white-collar workplaces looked like a lot like Don Draper's world on *Mad Men*. The offices with the windows and doors were occupied by men; the women were gathered in the secretarial pool in the center of the office, a sort of crude corral. Sexual access to them was considered a perk.

This might be why men in their sixties who are accused of behaving badly 30 years ago sometimes seem bewildered. They may feel they are being judged by contemporary standards for things they did under what they perceive as different rules. This is reflected in the data: According to a recent analysis by *The*

Economist, "younger respondents were more likely to think that a behavior crossed the line than their older peers were."

This does not absolve younger guys of their own bad behavior, nor is it reason to forgive the older men being accused. Still, it's important to talk more about these generational issues and how they color our thinking about the way we treat women.

How Do They Get Away With It?

Complicit assent. Think again about the *Access Hollywood* tape. What might have happened had Billy Bush, the show's host at the time, responded with, *Donald, that's disgusting—not to mention illegal!* Or if the other guys on the bus had said, *That's gross*. What if Harvey Weinstein's brother, Bob, had grabbed him by the shoulders and yelled, *Harvey, stop it! I will throw you out of the company if you continue!*

Sexual harassment persists because of three factors: the sense of entitlement that some men feel toward the women they work with; the presumption that women won't report it or fight back; and the presumed support—even tacit support in the form of not calling out bad behavior—of other men.

What we've seen recently is the second leg of the stool getting kicked out. There's been an outpouring of resistance from women. Women are speaking out, loudly, and not stopping.

What Can We Do?

Now it’s time to kick out the third leg. When men remain silent, it can be taken as a sign that we agree with the harasser, that we think the behavior is OK, and that we won’t intervene. Men are complicit in a culture that enables sexual harassment, so it is up to us to actively, volubly speak up and let the perpetrators know that we are not OK with what they do.

I’ll make one assertion here, which is backed by my experiences working with companies to promote gender equality over several decades: The overwhelming majority of men do not want to be jerks. We don’t want to make women uncomfortable and don’t want to say things that are offensive.

This puts a slightly more positive spin on the current male anxiety, which most assume is about being reported for harassment. But it also might be about the desire not to behave badly—and about not knowing exactly how to act.

We *can* act in a positive manner, however. Here’s one scenario I suspect is remarkably common:

Adeline is sitting in a meeting. She is the only woman in the room. Rob is in the meeting, too, and he makes a sexist comment. The room goes silent. Everyone’s attention is on Adeline: Is she going to do something, say something? *Oh, God, here she goes,* many of the other men are saying to themselves. Big eye roll. *She’s gonna call him out and make everyone feel bad*. And Adeline has to decide if she’s going to say something and make everyone miserable, or swallow it and stay miserable herself.

After the meeting, one of Adeline's colleagues, Fabrice, privately apologizes to her for Rob. "I'm really sorry about what he said in there," Fabrice says. "I didn't like that at all."

Fabrice thinks he's being supportive, but he's actually introducing another dilemma for Adeline. Does she nod politely and thank him? Or does she say, "Uh, where were you when I needed you?"

Men, what could you do differently? The obvious answer is that you could speak up, right then in the meeting, and say that you aren't comfortable with those kinds of statements. But typically we don't do that. Why not?

We're afraid that if we do, we'll be marginalized, kicked out of the men's club—that we'll become, in effect, "honorary women." Men know that doing the right thing sometimes carries costs, and most of us are worried about jeopardizing what we have. So we betray the women in the room, abandon our ethics, and slink away uncomfortably.

But think about that moment when Rob made his comment. I'm sure there were guys in the meeting who were looking down at their shoes, laughing uncomfortably, or shuffling the papers on the table. They didn't like it either but were too frightened to act.

Men, this is your chance. After the meeting, don't apologize to Adeline. Talk to one of the other guys who looked uneasy:

"Listen, Mateo, I hate it when Rob says things like that."

"So do I," says Mateo.

This is your opening: “The next time he does that, I’m going to say something. But as soon as I do, you have to jump right in and say that you don’t like it either. Can I count on you?”

Because here is what we know. It might be too scary for one guy to risk marginalization by speaking up, even though failing to do the right thing will make him ashamed later. But when two guys call out sexism, that opens a space for more men to chime in. And the behavior that makes women feel uncomfortable and alone might stop right there.

A global insurance company I consulted with developed informal “male allies” training, teaching men how to develop strategies to support one another. Critically, they were not being asked to “rescue” women; they were charged with challenging other men. The men developed several approaches, including supporting one another when a child was sick or a family issue arose. Soon the company’s male employees started talking more openly with one another about their experiences, their families, and their efforts to balance their lives. And after a year, the men reported higher levels of job satisfaction. Though it remains to be seen how these changes will affect sexual harassment at the company, the shared language and norms the men have developed will help them challenge one another and support men who speak out.

So, where do we go from here? After decades of accepting sexual harassment as the status quo, we have to take some of the weight off women’s shoulders. It’s simply not their responsibility

alone to talk about and enforce workplace equality. We must call out the sexist behaviors of other men because it's wrong and because it undermines women's confidence and effectiveness in the workplace.

This is what it means to be allies, men. To stand up together and do the right thing. We know how to do it, and we're good at it most of the time. Brotherhood, teamwork, and camaraderie are the essence of the fraternity, the foxhole, and the sports team. Now we have to learn how to come together at work—and on the right side of things.

Originally published in January 2018. Reprint BG1801

About the Contributors

SUSAN ASHFORD is the Michael & Susan Jandernoa Professor of Management and Organizations at the University of Michigan.

DOMINIC BARTON is the global managing partner of McKinsey & Company. He is a coauthor of *Talent Wins: The New Playbook for Putting People First* (Harvard Business Review Press, 2018).

NICHOLAS BLOOM is the William Eberle Professor of economics at Stanford University and a codirector of the Productivity, Innovation and Entrepreneurship program at the National Bureau of Economic Research.

JOSEPH L. BOWER is the Donald K. David Professor Emeritus at Harvard Business School. He is the coauthor (with Lynn S. Paine and Herman B. Leonard) of *Capitalism at Risk: Rethinking the Role of Business* (Harvard Business Review Press, 2011).

AARON K. CHATTERJI is an associate professor at Duke University's Fuqua School of Business and Sanford School of Public Policy.

J. YO-JUD CHENG is a doctoral candidate in the strategy unit at Harvard Business School.

SARAH CLIFFE is editorial director, *Harvard Business Review*.

THOMAS H. DAVENPORT is the President's Distinguished Professor of Information Technology and Management at Babson College, a research fellow at the MIT Initiative on the Digital Economy, and a senior adviser at Deloitte Analytics.

DEBBIE S. DOUGHERTY is a Professor in the Department of Communication at the University of Missouri and Editor Elect of the *Journal of Applied Communication Research*.

HEIDI K. GARDNER is a distinguished fellow at Harvard Law School and the author of *Smart Collaboration: How Professionals and Their Firms Succeed by Breaking Down Silos* (Harvard Business Review Press, 2017). Her research, teaching, and advisory work focus on leadership and collaboration in professional services firms.

BORIS GROYSBERG is the Richard P. Chapman Professor of Business Administration at Harvard Business School and a coauthor, with Michael Slind, of *Talk, Inc*. (Harvard Business Review Press, 2012). Twitter: @bgroysberg.

ALTON B. HARRIS is a law partner at Nixon Peabody. He has worked to promote gender equality in the workplace for more than 30 years. He is the coauthor (with Andrea S. Kramer) of *Breaking Through Bias: Communication Techniques*

for Women to Succeed at Work (Bibliomotion, 2016). Learn more at andieandal.com or on Twitter: @AndieandAl.

JAMES E. HEPPELMANN is the president and CEO of PTC, a leading maker of industrial software.

MARCO IANSITI is the David Sarnoff Professor of Business Administration at Harvard Business School, where he heads the Technology and Operations Management Unit and the Digital Initiative.

MICHAEL S. KIMMEL is the SUNY Distinguished Professor of Sociology and Gender Studies at Stony Brook University, where he founded the Center for the Study of Men and Masculinities, and the author of *Angry White Men* (Nation Books, 2013).

ANDREA S. KRAMER is a law partner at McDermott Will & Emery. She has worked to promote gender equality in the workplace for more than 30 years. She is the coauthor (with Alton B. Harris) of *Breaking Through Bias: Communication Techniques for Women to Succeed at Work* (Bibliomotion, 2016). Learn more at andieandal.com or on Twitter: @AndieandAl.

KARIM R. LAKHANI is a professor of business administration at Harvard Business School and the founding director of the Harvard Innovation Science Laboratory.

SUZANNE LEBSOCK is the Emeritus Board of Governors Professor of History at Rutgers University, where she cofounded the top-ranked graduate program in the history of women and gender. Lebsock's books include *A Murder in Virginia* (W.W. Norton, 2003) and *The Free Women of Petersburg* (W.W. Norton, 1984).

JEREMIAH LEE leads innovation for advisory services at Spencer Stuart. He is a cofounder of two culture-related businesses.

JAMES MANYIKA is the San Francisco–based director of the McKinsey Global Institute (MGI), the business and economics research arm of McKinsey & Company.

MARK MORTENSEN is an associate professor and the chair of the Organizational Behaviour Area at INSEAD. He researches, teaches, and consults on issues of collaboration, organizational design and new ways of working, and leadership.

LYNN S. PAINE is the John G. McLean Professor of Business Administration at HBS. She is the coauthor (with Joseph L. Bower and Herman B. Leonard) of *Capitalism at Risk: Rethinking the Role of Business* (Harvard Business Review Press, 2011).

GIANPIERO PETRIGLIERI is an associate professor of organizational behavior at INSEAD.

MICHAEL E. PORTER is a University Professor at Harvard, based at Harvard Business School.

JESSE PRICE is a leader in organizational culture services at Spencer Stuart. He is a cofounder of two culture-related businesses.

DAVID PYOTT is a philanthropist and the former CEO of Allergan.

RAJEEV RONANKI is a principal at Deloitte Consulting, where he leads the cognitive computing and health care innovation practices. Some of the companies mentioned in his article are Deloitte clients.

RAFFAELLA SADUN is the Thomas S. Murphy Associate Professor of Business Administration at Harvard Business School.

MICHAEL W. TOFFEL is the Senator John Heinz Professor of Environmental Management at Harvard Business School.

JOHN VAN REENEN is the Gordon Y. Billard Professor in MIT's department of economics and its Sloan School of Management.

MAXINE WILLIAMS is Facebook's global director of diversity.

JOAN C. WILLIAMS is Distinguished Professor of Law and Founding Director of the Center of WorkLife Law at the University of California, Hastings College of the Law. Her newest book is *White Working Class: Overcoming Class Cluelessness in America* (Harvard Business Review Press, 2017).

SARAH KEOHANE WILLIAMSON is CEO of FCLT Global.

AMY WRZESNIEWSKI is a professor of organizational behavior at the Yale School of Management.

Index

Invaluable insights always at your fingertips

With an All-Access subscription to *Harvard Business Review*, you'll get so much more than a magazine.

Exclusive online content and tools you can put to use today

My Library, your personal workspace for sharing, saving, and organizing HBR.org articles and tools

Unlimited access to more than 4,000 articles in the *Harvard Business Review* archive

Subscribe today at hbr.org/subnow

The most important management ideas all in one place.

We hope you enjoyed this book from *Harvard Business Review*. For the best ideas HBR has to offer turn to HBR's 10 Must Reads Boxed Set. From books on leadership and strategy to managing yourself and others, this 6-book collection delivers articles on the most essential business topics to help you succeed.

HBR's 10 Must Reads Series

The definitive collection of ideas and best practices on our most sought-after topics from the best minds in business.

- Change Management
- Collaboration
- Communication
- Emotional Intelligence
- Innovation
- Leadership
- Making Smart Decisions
- Managing Across Cultures
- Managing People
- Managing Yourself
- Strategic Marketing
- Strategy
- Teams
- The Essentials

hbr.org/mustreads

Buy for your team, clients, or event.
Visit hbr.org/bulksales for quantity discount rates.

Harvard Business Review Press

HBR'S 10 MUST READS

BONUS ARTICLE

"How CEOs Manage Time"

By Michael E. Porter and Nitin Nohria

The definitive management ideas of the year from **Harvard Business Review.**

2020

HBR's 10 Must Reads series is the definitive collection of ideas and best practices for aspiring and experienced leaders alike. These books offer essential reading selected from the pages of *Harvard Business Review* on topics critical to the success of every manager.

Titles include:

HBR's 10 Must Reads 2015
HBR's 10 Must Reads 2016
HBR's 10 Must Reads 2017
HBR's 10 Must Reads 2018
HBR's 10 Must Reads 2019
HBR's 10 Must Reads 2020
HBR's 10 Must Reads for CEOs
HBR's 10 Must Reads for New Managers
HBR's 10 Must Reads on AI, Analytics, and the New Machine Age
HBR's 10 Must Reads on Business Model Innovation
HBR's 10 Must Reads on Change Management
HBR's 10 Must Reads on Collaboration
HBR's 10 Must Reads on Communication
HBR's 10 Must Reads on Diversity
HBR's 10 Must Reads on Emotional Intelligence
HBR's 10 Must Reads on Entrepreneurship and Startups
HBR's 10 Must Reads on Innovation

HBR's 10 Must Reads on Leadership
HBR's 10 Must Reads on Leadership for Healthcare
HBR's 10 Must Reads on Leadership Lessons from Sports
HBR's 10 Must Reads on Making Smart Decisions
HBR's 10 Must Reads on Managing Across Cultures
HBR's 10 Must Reads on Managing People
HBR's 10 Must Reads on Managing Yourself
HBR's 10 Must Reads on Mental Toughness
HBR's 10 Must Reads on Negotiation
HBR's 10 Must Reads on Nonprofits and the Social Sectors
HBR's 10 Must Reads on Reinventing HR
HBR's 10 Must Reads on Sales
HBR's 10 Must Reads on Strategic Marketing
HBR's 10 Must Reads on Strategy
HBR's 10 Must Reads on Strategy for Healthcare
HBR's 10 Must Reads on Teams
HBR's 10 Must Reads on Women and Leadership
HBR's 10 Must Reads: The Essentials

HBR'S 10 MUST READS

The definitive
management ideas
of the year from
Harvard Business Review.
2020

HARVARD BUSINESS REVIEW PRESS

Boston, Massachusetts

HBR Press Quantity Sales Discounts

Harvard Business Review Press titles are available at significant quantity discounts when purchased in bulk for client gifts, sales promotions, and premiums. Special editions, including books with corporate logos, customized covers, and letters from the company or CEO printed in the front matter, as well as excerpts of existing books, can also be created in large quantities for special needs.

For details and discount information for both print and ebook formats, contact booksales@harvardbusiness.org, tel. 800-988-0886, or www.hbr.org/bulksales.

Copyright 2020 Harvard Business Publishing Corporation

All rights reserved

No part of this publication may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form, or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the prior permission of the publisher. Requests for permission should be directed to permissions@harvardbusiness.org, or mailed to Permissions, Harvard Business School Publishing, 60 Harvard Way, Boston, Massachusetts 02163.

First eBook Edition: Oct 2019

ISBN: 978-1-63369-812-3
eISBN: 978-1-63369-813-0

Contents

Editors' Note

Each year as we pare twelve months' worth of HBR articles down to a handful of the very best, certain trends and themes emerge. Sometimes the economy, politics, and technology loom largest. Other times new twists on the basics of leadership, strategy, and marketing dominate. This year's choices are united not by a trend or a theme but by a *feeling*: the surprise we experience when some long-held truth is gently challenged and is revealed to be different or more complex than we had thought. Educators and psychologists know that novelty reinforces understanding and learning. Indeed, the sense of surprise that distinguishes the pieces in this collection makes them stay with us, tugging at loose threads in our minds, helping us see connections as we seek to improve our organizations and ourselves. It spurs us out of complacency and encourages the mindset needed to learn, grow, and innovate. Embrace that feeling while you read and as you and your business prepare for coming challenges.

Some professionals are trained to ask questions: Think of litigators, journalists, and doctors. But few executives regard questioning as a skill to be honed. That's a missed opportunity, say Alison Wood Brooks and Leslie K. John. Thoughtful inquiry and the conversational exchange of ideas can yield a kind of magic, a whole that is greater than the sum of its parts. In "**The Surprising Power of Questions**," the authors describe several

carefully researched techniques that will help you boost your learning, persuade others, and negotiate more effectively.

The traditional analytical tools of strategy may be well suited to understanding an existing business context, but they're of little value when you need to reinvent your business. To generate a breakthrough strategy, Adam Brandenburger suggests building one with tools explicitly designed to foster creativity. "**Strategy Needs Creativity**" details approaches for kindling a spark of intuition, making a connection between disparate ways of thinking, or taking a leap into the unexpected that can lead you to a game-changing way of doing business.

"Women lack the desire and ability to negotiate"; "Women are more committed to family than men are." Statements like these are often proffered to explain why women have failed to gain parity with men in the workplace, but Catherine H. Tinsley and Robin J. Ely fell them with a few deft strokes. "Science, by and large, does not actually support these claims," they write. "The sexes are far more similar in their inclinations, attitudes, and skills than popular opinion would have us believe." The authors show that managers who are advancing gender equity in their firms take an inquisitive approach: They seek an evidence-based understanding of how women experience the workplace, and then create conditions that increase their prospects for success. "**What Most People Get Wrong About Men and Women**" is a clarion call for rejecting the script that

encourages women to act more like men and instead fixing the things that undermine women and reinforce gender stereotypes.

Artificial intelligence is becoming good at many "human" jobs—diagnosing disease, translating languages, providing customer service. And it's improving fast, raising reasonable fears that AI will ultimately replace human workers throughout the economy. Accenture technology leaders H. James Wilson and Paul R. Daugherty argue that that's not the inevitable, or even the most likely, outcome. In "**Collaborative Intelligence: Humans and AI Are Joining Forces**," they show that AI has the most significant impact—and companies see the biggest performance gains—when people and smart machines work together, enhancing one another's strengths. Organizations that use machines merely to displace workers, they say, will miss the full potential of AI.

Stitch Fix demonstrates human and machine collaboration in action. The company has a simple business model: It sends you clothing and accessories it thinks you'll like; you keep the items you want and send the others back. But behind the curtain is a relentlessly data-driven organization built on the belief that a good person plus a good algorithm is better than either the best person or the best algorithm alone. In "**Stitch Fix's CEO on Selling Personal Style to the Mass Market**," company founder and chief executive Katrina Lake describes overcoming skeptics (one of her business school professors called her idea

an “inventory nightmare”) and surmounting the challenges of raising capital for a clothing start-up in the male-dominated VC field.

Is “**Strategy for Start-Ups**” the beginning of a new paradigm, or is it entrepreneurial heresy? HBR’s most divisive article of the year details how entrepreneurs often run with the first plausible strategy they identify in their haste to get to market. As a result they lose out to second or even third movers with superior strategies. Having worked with and studied hundreds of start-ups over the past 20 years, Joshua Gans, Erin L. Scott, and Scott Stern have developed a framework that helps founders take a practical, clarifying approach to the critical choices they face. The authors delineate four go-to-market strategies for entrepreneurs to consider as they move from idea to launch. Each option offers a distinct way for the venture to create and capture value.

Agile innovation teams are small, entrepreneurial groups designed to stay close to customers and adapt quickly to changing conditions. When implemented correctly, they almost always result in greater productivity, better morale, faster time to market, higher quality, and lower risk than traditional approaches can achieve. In “**Agile at Scale**,” Darrell K. Rigby, Jeff Sutherland, and Andy Noble explore how your company can go from a handful of agile teams to hundreds. Making agile the dominant way you operate, they say, means committing all the way to the top: Leaders should adopt agile values, create a

taxonomy of opportunities to set priorities, and break the journey of transformation into small steps.

Conventional wisdom holds that the more contact an operation has with its customers, the less efficiently it will run. But when customers are walled off, they are unlikely to fully understand and appreciate the work going on behind the scenes. "**Operational Transparency**" advocates for the deliberate design of windows into and out of an organization's processes so that customers can recognize the value being added. Take open kitchens: Research shows that when diners can see who's making their food, their satisfaction increases—and it's even greater if the chef can see the diners. Ryan W. Buell describes how managers can bring this sort of transparency to their companies, exploring what to reveal, when to reveal it, and how to avoid going too far.

What does it take for companies to both do well and do good? Many corporations are seeking to dial down their single-minded pursuit of financial gain and pay closer attention to their impact on the environment and society—but the business ecosystem is still motivated above all by shareholder wealth. In "**The Dual-Purpose Playbook**," Julie Battilana, Anne-Claire Pache, Metin Sengul, and Marissa Kimsey look at how companies can find a balance. Examining dual-purpose companies around the globe, they find that successful ones build a commitment to both economic and social value into their core organizational activities. The authors outline four key management practices,

which range from setting and monitoring dual goals to hiring and socializing employees to embrace them.

Strategy guru Michael E. Porter and Harvard Business School dean Nitin Nohria teamed up for 12 years to collect 60,000 hours' worth of data from 27 chief executives, all to better understand what their days consist of. Although CEOs have tremendous resources at their disposal, time remains an area of acute scarcity. "**How CEOs Manage Time**," a fine-grained, first-of-a-kind study, reveals similarities in how CEOs structure their schedules (they all attend a lot of meetings) along with differences (some dedicate far more face time to investors and customers than others do). It sheds light on the crucial trade-offs executives must make and describes how any leader can manage his or her calendar more effectively.

It's often called the silver tsunami: In many countries the population is aging rapidly. In the United States alone about 10,000 people turn 65 each day—and one in five Americans will be 65 or older by 2030. This societal shift will affect every aspect of business, but Paul Irving finds that many corporate leaders have not yet considered its effects. And those who have, he says, typically foresee a looming crisis and miss the potential contributions that older adults—healthier and more active than their predecessors—can make as both workers and consumers. The final piece in this volume, "**When No One Retires**," helps companies develop a "longevity strategy" for fostering a vibrant multigenerational workforce.

These standout pieces of the year explore some of the most compelling and important developments in business today. Did anything you just read in these descriptions surprise you? We hope so. Use that feeling to view yourself and your business through a new lens as you seek to improve and grow in the coming year.

—The Editors

The Surprising Power of Questions

by Alison Wood Brooks and Leslie K. John

MUCH OF AN EXECUTIVE'S WORKDAY is spent asking others for information—requesting status updates from a team leader, for example, or questioning a counterpart in a tense negotiation. Yet unlike professionals such as litigators, journalists, and doctors, who are taught how to ask questions as an essential part of their training, few executives think of questioning as a skill that can be honed—or consider how their own answers to questions could make conversations more productive.

That's a missed opportunity. Questioning is a uniquely powerful tool for unlocking value in organizations: It spurs learning and the exchange of ideas, it fuels innovation and performance improvement, it builds rapport and trust among team members. And it can mitigate business risk by uncovering unforeseen pitfalls and hazards.

For some people, questioning comes easily. Their natural inquisitiveness, emotional intelligence, and ability to read people put the ideal question on the tip of their tongue. But most of us

don't ask enough questions, nor do we pose our inquiries in an optimal way.

The good news is that by asking questions, we naturally improve our emotional intelligence, which in turn makes us better questioners—a virtuous cycle. In this article, we draw on insights from behavioral science research to explore how the way we frame questions and choose to answer our counterparts can influence the outcome of conversations. We offer guidance for choosing the best type, tone, sequence, and framing of questions and for deciding what and how much information to share to reap the most benefit from our interactions, not just for ourselves but for our organizations.

Don't Ask, Don't Get

"Be a good listener," Dale Carnegie advised in his 1936 classic *How to Win Friends and Influence People*. "Ask questions the other person will enjoy answering." More than 80 years later, most people still fail to heed Carnegie's sage advice. When one of us (Alison) began studying conversations at Harvard Business School several years ago, she quickly arrived at a foundational insight: People don't ask enough questions. In fact, among the most common complaints people make after having a conversation, such as an interview, a first date, or a work meeting, is "I wish [s/he] had asked me more questions" and "I can't believe [s/he] didn't ask me any questions."

Why do so many of us hold back? There are many reasons. People may be egocentric—eager to impress others with their own thoughts, stories, and ideas (and not even think to ask questions). Perhaps they are apathetic—they don't care enough to ask, or they anticipate being bored by the answers they'd hear. They may be overconfident in their own knowledge and think they already know the answers (which sometimes they do, but usually not). Or perhaps they worry that they'll ask the wrong question and be viewed as rude or incompetent. But the biggest inhibitor, in our opinion, is that most people just don't understand how beneficial good questioning can be. If they did, they would end far fewer sentences with a period—and more with a question mark.

Dating back to the 1970s, research suggests that people have conversations to accomplish some combination of two major goals: information exchange (learning) and impression management (liking). Recent research shows that asking questions achieves both. Alison and Harvard colleagues Karen Huang, Michael Yeomans, Julia Minson, and Francesca Gino scrutinized thousands of natural conversations among participants who were getting to know each other, either in online chats or on in-person speed dates. The researchers told some people to ask many questions (at least nine in fifteen minutes) and others to ask very few (no more than four in fifteen minutes). In the online chats, the people who were randomly assigned to ask many questions were better liked by their

conversation partners and learned more about their partners' interests. For example, when quizzed about their partners' preferences for activities such as reading, cooking, and exercising, high question askers were more likely to be able to guess correctly. Among the speed daters, people were more willing to go on a second date with partners who asked more questions. In fact, asking just one more question on each date meant that participants persuaded one additional person (over the course of 20 dates) to go out with them again.

Idea in Brief

The Problem

Some professionals such as litigators, journalists, and even doctors are taught to ask questions as part of their training. But few executives think about questioning as a skill that can be honed. That's a missed opportunity.

The Opportunity

Questioning is a powerful tool for unlocking value in companies: It spurs learning and the exchange of ideas, it fuels innovation and better performance, it builds trust among team members. And it can mitigate business risk by uncovering unforeseen pitfalls and hazards.

The Approach

Several techniques can enhance the power and efficacy of queries: Favor follow-up questions, know when to keep questions open-ended, get the sequence right, use the right tone, and pay attention to group dynamics.

Questions are such powerful tools that they can be beneficial—perhaps particularly so—in circumstances when question asking

goes against social norms. For instance, prevailing norms tell us that job candidates are expected to answer questions during interviews. But research by Dan Cable, at the London Business School, and Virginia Kay, at the University of North Carolina, suggests that most people excessively self-promote during job interviews. And when interviewees focus on selling themselves, they are likely to forget to ask questions—about the interviewer, the organization, the work—that would make the interviewer feel more engaged and more apt to view the candidate favorably and could help the candidate predict whether the job would provide satisfying work. For job candidates, asking questions such as "What am I not asking you that I should?" can signal competence, build rapport, and unlock key pieces of information about the position.

Most people don't grasp that asking a lot of questions unlocks learning and improves interpersonal bonding. In Alison's studies, for example, though people could accurately recall how many questions had been asked in their conversations, they didn't intuit the link between questions and liking. Across four studies, in which participants were engaged in conversations themselves or read transcripts of others' conversations, people tended not to realize that question asking would influence—or had influenced—the level of amity between the conversationalists.

The New Socratic Method

The first step in becoming a better questioner is simply to ask more questions. Of course, the sheer number of questions is not the only factor that influences the quality of a conversation: The type, tone, sequence, and framing also matter.

In our teaching at Harvard Business School, we run an exercise in which we instruct pairs of students to have a conversation. Some students are told to ask as few questions as possible, and some are instructed to ask as many as possible. Among the low-low pairs (both students ask a minimum of questions), participants generally report that the experience is a bit like children engaging in parallel play: They exchange statements but struggle to initiate an interactive, enjoyable, or productive dialogue. The high-high pairs find that too many questions can also create a stilted dynamic. However, the high-low pairs' experiences are mixed. Sometimes the question asker learns a lot about her partner, the answerer feels heard, and both come away feeling profoundly closer. Other times, one of the participants may feel uncomfortable in his role or unsure about how much to share, and the conversation can feel like an interrogation.

Our research suggests several approaches that can enhance the power and efficacy of queries. The best approach for a given situation depends on the goals of the conversationalists—specifically, whether the discussion is cooperative (for example, the duo is trying to build a relationship or accomplish a task together) or competitive (the parties seek to uncover sensitive information from each other or serve their own interests), or

some combination of both. (See the sidebar "Conversational Goals Matter.") Consider the following tactics.

Favor follow-up questions

Not all questions are created equal. Alison's research, using human coding and machine learning, revealed four types of questions: introductory questions ("How are you?"), mirror questions ("I'm fine. How are you?"), full-switch questions (ones that change the topic entirely), and follow-up questions (ones that solicit more information). Although each type is abundant in natural conversation, follow-up questions seem to have special power. They signal to your conversation partner that you are listening, care, and want to know more. People interacting with a partner who asks lots of follow-up questions tend to feel respected and heard.

An unexpected benefit of follow-up questions is that they don't require much thought or preparation—indeed, they seem to come naturally to interlocutors. In Alison's studies, the people who were told to ask more questions used more follow-up questions than any other type without being instructed to do so.

Know when to keep questions open-ended

No one likes to feel interrogated—and some types of questions can force answerers into a yes-or-no corner. Open-ended questions can counteract that effect and thus can be particularly useful in uncovering information or learning something new. Indeed, they are wellsprings of innovation—which is often the

result of finding the hidden, unexpected answer that no one has thought of before.

Conversational goals matter

Conversations fall along a continuum from purely competitive to purely cooperative. For example, discussions about the allocation of scarce resources tend to be competitive; those between friends and colleagues are generally cooperative; and others, such as managers' check-ins with employees, are mixed—supportive but also providing feedback and communicating expectations. Here are some challenges that commonly arise when asking and answering questions and tactics for handling them.

Competitive conversations

When you're asking

Challenge

Your conversational partner is reluctant to share information and may even lie.

Tactics

Ask direct or "yes or no" questions to avoid evasive answers.

Ask detailed follow-up questions (even if they're redundant) to pry out more information.

Frame tough questions using pessimistic assumptions ("We've experienced some headwinds in sales, right?") to reduce the likelihood that respondents will lie.

Ask the most sensitive question first. Subsequent questions will feel less intrusive, making your partner more forthcoming.

When you're answering

Challenge

Answering questions could put you at a strategic disadvantage.

Tactics

Prepare: Think in advance about the information you want to keep private to avoid answering impulsively.

Dodge the issue by answering a similar question you'd prefer to have been asked.

Deflect and gain control of the conversation by posing a question in return.

Consider when to share negative information—rather than refusing to answer—to build trust.

Cooperative conversations

When you're asking

Challenge

Friendly colleagues may shy away from conflict or hesitate to share bad news.

Tactics

Ask open-ended questions ("If you were to play devil's advocate, what would you say?") to draw out negative feedback.

Begin with the least-sensitive questions to build rapport, and escalate slowly.

As in competitive contexts, frame tough questions using negative assumptions.

When you're answering

Challenge

We forget to focus on making conversations productive and delightful, or we speak too freely.

Tactics

Avoid droning on and on. Use energy, humor, and storytelling to engage your partners.

Avoid talking too much about yourself, and remember to ask questions of others.

Deflect tough questions by answering with another question or a joke.

A wealth of research in survey design has shown the dangers of narrowing respondents' options. For example, "closed" questions can introduce bias and manipulation. In one study, in which parents were asked what they deemed "the most important thing for children to prepare them in life," about 60% of them chose "to think for themselves" from a list of response options. However, when the same question was asked in an open-ended format, only about 5% of parents spontaneously came up with an answer along those lines.

Of course, open-ended questions aren't always optimal. For example, if you are in a tense negotiation or are dealing with people who tend to keep their cards close to their chest, open-ended questions can leave too much wiggle room, inviting them to dodge or lie by omission. In such situations, closed questions work better, especially if they are framed correctly. For example, research by Julia Minson, the University of Utah's Eric VanEpps, Georgetown's Jeremy Yip, and Wharton's Maurice Schweitzer indicates that people are less likely to lie if questioners make pessimistic assumptions ("This business will need some new equipment soon, correct?") rather than optimistic ones ("The equipment is in good working order, right?").

Sometimes the information you wish to ascertain is so sensitive that direct questions won't work, no matter how thoughtfully they are framed. In these situations, a survey tactic can aid discovery. In research Leslie conducted with Alessandro

Acquisti and George Loewenstein of Carnegie Mellon University, she found that people were more forthcoming when requests for sensitive information were couched within another task—in the study's case, rating the ethicality of antisocial behaviors such as cheating on one's tax return or letting a drunk friend drive home. Participants were asked to rate the ethicality using one scale if they had engaged in a particular behavior and another scale if they hadn't—thus revealing which antisocial acts they themselves had engaged in. Although this tactic may sometimes prove useful at an organizational level—we can imagine that managers might administer a survey rather than ask workers directly about sensitive information such as salary expectations —we counsel restraint in using it. If people feel that you are trying to trick them into revealing something, they may lose trust in you, decreasing the likelihood that they'll share information in the future and potentially eroding workplace relationships.

Get the sequence right

The optimal order of your questions depends on the circumstances. During tense encounters, asking tough questions first, even if it feels socially awkward to do so, can make your conversational partner more willing to open up. Leslie and her coauthors found that people are more willing to reveal sensitive information when questions are asked in a decreasing order of intrusiveness. When a question asker begins with a highly sensitive question—such as "Have you ever had a fantasy of

doing something terrible to someone?"—subsequent questions, such as "Have you ever called in sick to work when you were perfectly healthy?" feel, by comparison, less intrusive, and thus we tend to be more forthcoming. Of course, if the first question is *too* sensitive, you run the risk of offending your counterpart. So it's a delicate balance, to be sure.

If the goal is to build relationships, the opposite approach—opening with less sensitive questions and escalating slowly—seems to be most effective. In a classic set of studies (the results of which went viral following a write-up in the "Modern Love" column of the *New York Times*), psychologist Arthur Aron recruited strangers to come to the lab, paired them up, and gave them a list of questions. They were told to work their way through the list, starting with relatively shallow inquiries and progressing to more self-revelatory ones, such as "What is your biggest regret?" Pairs in the control group were asked simply to interact with each other. The pairs who followed the prescribed structure liked each other more than the control pairs. This effect is so strong that it has been formalized in a task called "the relationship closeness induction," a tool used by researchers to build a sense of connection among experiment participants.

Good interlocutors also understand that questions asked previously in a conversation can influence future queries. For example, Norbert Schwarz, of the University of Southern California, and his coauthors found that when the question "How satisfied are you with your life?" is followed by the

question “How satisfied are you with your marriage?” the answers were highly correlated: Respondents who reported being satisfied with their life also said they were satisfied with their marriage. When asked the questions in this order, people implicitly interpreted that life satisfaction “ought to be” closely tied to marriage. However, when the same questions were asked in the opposite order, the answers were less closely correlated.

The Power of Questions in Sales

THERE ARE FEW BUSINESS SETTINGS in which asking questions is more important than sales. A recent study of more than 500,000 business-to-business sales conversations—over the phone and via online platforms—by tech company Gong.io reveals that top-performing salespeople ask questions differently than their peers.

Consistent with past research, the data shows a strong connection between the number of questions a salesperson asks and his or her sales conversion rate (in terms of both securing the next meeting and eventually closing the deal). This is true even after controlling for the gender of the salesperson and the call type (demo, proposal, negotiation, and so on). However, there is a point of diminishing returns. Conversion rates start to drop off after about 14 questions, with 11 to 14 being the optimal range.

The data also shows that top-performing salespeople tend to scatter questions throughout the sales call, which makes it feel more like a conversation than an interrogation. Lower performers, in contrast, frontload questions in the first half of the sales call, as if they’re making their way through a to-do list.

Just as important, top salespeople listen more and speak less than their counterparts overall. Taken together, the data from Gong.io affirms what

great salespeople intuitively understand: When sellers ask questions rather than just make their pitch, they close more deals.

Use the right tone

People are more forthcoming when you ask questions in a casual way, rather than in a buttoned-up, official tone. In one of Leslie's studies, participants were posed a series of sensitive questions in an online survey. For one group of participants, the website's user interface looked fun and frivolous; for another group, the site looked official. (The control group was presented with a neutral-looking site.) Participants were about twice as likely to reveal sensitive information on the casual-looking site than on the others.

People also tend to be more forthcoming when given an escape hatch or "out" in a conversation. For example, if they are told that they can change their answers at any point, they tend to open up more—even though they rarely end up making changes. This might explain why teams and groups find brainstorming sessions so productive. In a whiteboard setting, where anything can be erased and judgment is suspended, people are more likely to answer questions honestly and say things they otherwise might not. Of course, there will be times when an off-the-cuff approach is inappropriate. But in general, an overly formal tone is likely to inhibit people's willingness to share information.

Pay attention to group dynamics

Conversational dynamics can change profoundly depending on whether you're chatting one-on-one with someone or talking in a group. Not only is the willingness to answer questions affected simply by the presence of others, but members of a group tend to follow one another's lead. In one set of studies, Leslie and her coauthors asked participants a series of sensitive questions, including ones about finances ("Have you ever bounced a check?") and sex ("While an adult, have you ever felt sexual desire for a minor?"). Participants were told either that most others in the study were willing to reveal stigmatizing answers or that they were unwilling to do so. Participants who were told that others had been forthcoming were 27% likelier to reveal sensitive answers than those who were told that others had been reticent. In a meeting or group setting, it takes only a few closed-off people for questions to lose their probing power. The opposite is true, too. As soon as one person starts to open up, the rest of the group is likely to follow suit.

Group dynamics can also affect how a question asker is perceived. Alison's research reveals that participants in a conversation enjoy being asked questions and tend to like the people asking questions more than those who answer them. But when third-party observers watch the same conversation unfold, they prefer the person who answers questions. This makes sense: People who mostly ask questions tend to disclose very little about themselves or their thoughts. To those listening to a conversation, question askers may come across as defensive,

evasive, or invisible, while those answering seem more fascinating, present, or memorable.

The Best Response

A conversation is a dance that requires partners to be in sync—it's a mutual push-and-pull that unfolds over time. Just as the way we ask questions can facilitate trust and the sharing of information—so, too, can the way we answer them.

Answering questions requires making a choice about where to fall on a continuum between privacy and transparency. Should we answer the question? If we answer, how forthcoming should we be? What should we do when asked a question that, if answered truthfully, might reveal a less-than-glamorous fact or put us in a disadvantaged strategic position? Each end of the spectrum—fully opaque and fully transparent—has benefits and pitfalls. Keeping information private can make us feel free to experiment and learn. In negotiations, withholding sensitive information (such as the fact that your alternatives are weak) can help you secure better outcomes. At the same time, transparency is an essential part of forging meaningful connections. Even in a negotiation context, transparency can lead to value-creating deals; by sharing information, participants can identify elements that are relatively unimportant to one party but important to the other—the foundation of a win-win outcome.

And keeping secrets has costs. Research by Julie Lane and Daniel Wegner, of the University of Virginia, suggests that

concealing secrets during social interactions leads to the intrusive recurrence of secret thoughts, while research by Columbia's Michael Slepian, Jinseok Chun, and Malia Mason shows that keeping secrets—even outside of social interactions—depletes us cognitively, interferes with our ability to concentrate and remember things, and even harms long-term health and well-being.

In an organizational context, people too often err on the side of privacy—and underappreciate the benefits of transparency. How often do we realize that we could have truly bonded with a colleague only after he or she has moved on to a new company? Why are better deals often uncovered after the ink has dried, the tension has broken, and negotiators begin to chat freely?

To maximize the benefits of answering questions—and minimize the risks—it's important to decide before a conversation begins what information you want to share and what you want to keep private.

Deciding what to share

There is no rule of thumb for how much—or what type—of information you should disclose. Indeed, transparency is such a powerful bonding agent that sometimes it doesn't matter what is revealed—even information that reflects poorly on us can draw our conversational partners closer. In research Leslie conducted with HBS collaborators Kate Barasz and Michael Norton, she found that most people assume that it would be less damaging to refuse to answer a question that would reveal negative

information—for example, “Have you ever been reprimanded at work?”—than to answer affirmatively. But this intuition is wrong. When they asked people to take the perspective of a recruiter and choose between two candidates (equivalent except for how they responded to this question), nearly 90% preferred the candidate who “came clean” and answered the question. Before a conversation takes place, think carefully about whether refusing to answer tough questions would do more harm than good.

Deciding what to keep private

Of course, at times you and your organization would be better served by keeping your cards close to your chest. In our negotiation classes, we teach strategies for handling hard questions without lying. Dodging, or answering a question you *wish* you had been asked, can be effective not only in helping you protect information you’d rather keep private but also in building a good rapport with your conversational partner, especially if you speak eloquently. In a study led by Todd Rogers, of Harvard’s Kennedy School, participants were shown clips of political candidates responding to questions by either answering them or dodging them. Eloquent dodgers were liked more than ineloquent answerers, but only when their dodges went undetected. Another effective strategy is deflecting, or answering a probing question with another question or a joke. Answerers can use this approach to lead the conversation in a different direction.

"Question everything," Albert Einstein famously said. Personal creativity and organizational innovation rely on a willingness to seek out novel information. Questions and thoughtful answers foster smoother and more-effective interactions; they strengthen rapport and trust and lead groups toward discovery. All this we have documented in our research. But we believe questions and answers have a power that goes far beyond matters of performance. The wellspring of all questions is wonder and curiosity and a capacity for delight. We pose and respond to queries in the belief that the magic of a conversation will produce a whole that is greater than the sum of its parts. Sustained personal engagement and motivation—in our lives as well as our work—require that we are always mindful of the transformative joy of asking and answering questions.

Originally published in May–June 2018. Reprint R1803C

Strategy Needs Creativity

by Adam Brandenburger

I'VE NOTICED THAT BUSINESS SCHOOL STUDENTS often feel frustrated when they're taught strategy. There's a gap between what they learn and what they'd like to learn. Strategy professors (including me) typically teach students to think about strategy problems by introducing them to rigorous analytical tools—assessing the five forces, drawing a value net, plotting competitive positions. The students know that the tools are essential, and they dutifully learn how to use them. But they also realize that the tools are better suited to understanding an existing business context than to dreaming up ways to reshape it. Game-changing strategies, they know, are born of creative thinking: a spark of intuition, a connection between different ways of thinking, a leap into the unexpected.

They're right to feel this way—which is not to say that we should abandon the many powerful analytical tools we've developed over the years. We'll always need them to understand competitive landscapes and to assess how

companies can best deploy their resources and competencies there. But we who devote our professional lives to thinking about strategy need to acknowledge that just giving people those tools will not help them break with conventional ways of thinking. If we want to teach students—and executives—how to generate groundbreaking strategies, we must give them tools explicitly designed to foster creativity.

A number of such tools already exist, often in practitioner-friendly forms. In "How Strategists Really Think: Tapping the Power of Analogy" (HBR, April 2005), Giovanni Gavetti and Jan W. Rivkin write compellingly about using analogies to come up with new business models. Charles Duhigg talks in his book *Smarter Faster Better* about introducing carefully chosen creative "disturbances" into work processes to spur new thinking. Youngme Moon, in "Break Free from the Product Life Cycle" (HBR, May 2005), suggests redefining products by boldly limiting—rather than augmenting—the features offered.

What these approaches have in common is the goal of moving strategy past the insights delivered by analytic tools (which are close at hand) and into territory that's further afield, or—to use a bit of academic jargon—*cognitively distant.* They take their inspiration more from how our thought processes work than from how industries or business models are structured. For that reason they can help strategists make the creative leap beyond what already exists to invent a genuinely

new way of doing business. Simply waiting for inspiration to strike is not the answer.

In this article I explore four approaches to building a breakthrough strategy:

- **Contrast.** The strategist should identify—and challenge—the assumptions undergirding the company's or the industry's status quo. This is the most direct and often the most powerful way to reinvent a business.
- **Combination.** Steve Jobs famously said that creativity is "just connecting things"; many smart business moves come from linking products or services that seem independent from or even in tension with one another.
- **Constraint.** A good strategist looks at an organization's limitations and considers how they might actually become strengths.
- **Context.** If you reflect on how a problem similar to yours was solved in an entirely different context, surprising insights may emerge. (I wrote about these ideas more academically in "Where Do Great Strategies Really Come From?" *Strategy Science*, December 2017.)

These approaches aren't exhaustive—or even entirely distinct from one another—but I've found that they help people explore a wide range of possibilities.

Idea in Brief

The Problem

The field of strategy overfocuses on analytic rigor and underfocuses on creativity.

Why It Matters

Analytic tools are good at helping strategists develop business ideas that are close at hand—but less good at discovering transformative strategies.

In Practice

The wise strategist can work with four creativity-enhancing tools: contrast, combination, constraint, and context.

Contrast: What Pieces of Conventional Wisdom Are Ripe for Contradiction?

To create a strategy built on contrast, first identify the assumptions implicit in existing strategies. Elon Musk seems to have a knack for this approach. He and the other creators of PayPal took a widely held but untested assumption about banking—that transferring money online was feasible and safe between institutions but not between individuals—and disproved it. With SpaceX he is attempting to overturn major assumptions about space travel: that it must occur on a fixed schedule, be paid for by the public, and use onetime rockets. He may be on track toward a privately funded, on-demand business that reuses rockets.

It's best to be precise—even literal—when naming such assumptions. Consider the video rental industry in 2000. Blockbuster ruled the industry, and the assumptions beneath its model seemed self-evident: People pick up videos at a retail location close to home. Inventory must be limited because new videos are expensive. Since the demand for them is high, customers must be charged for late returns. (It was basically a public-library model.) But Netflix put those assumptions under a microscope. Why is a physical location necessary? Mailing out videos would be cheaper and more convenient. Is there a way around the high fees for new releases? If the studios were open to a revenue-sharing agreement, both parties could benefit. Those two changes allowed Netflix to carry lots more movies, offer long rental periods, do away with late fees—and remake an industry.

Most of the time, strategy from contrast may look less revolutionary than Netflix (which remade itself again by streaming videos and becoming a content creator) or SpaceX (should it succeed). Any organization can ask whether it might usefully flip the order in which it performs activities, for example. The traditional model in retail is to start with a flagship store (usually in a city center) and add satellites (in suburban locations). Now consider pop-up stores: In some cases they conform to the old model—they are like mini-satellites; but in others the pop-up comes first, and if that's successful, a

larger footprint is added. The Soho area of New York City has become a testing ground for this strategy.

Another approach is to consider shaking up the value chain, which in any industry is conventionally oriented in a particular way, with some players acting as suppliers and others as customers. Inverting the value chain may yield new business models. In the charitable sector, for example, donors have been seen as suppliers of financial resources. DonorsChoose.org is a model that treats them more like customers. The organization puts up a “storefront” of requests posted by schoolteachers around the United States who are looking for materials for their (often underresourced) classrooms. Donors can choose which requests to respond to and receive photos of the schoolwork that their money has supported. In effect, they are buying the satisfaction of seeing a particular classroom before and after.

In some industries the status quo has dictated highly bundled, expensive products or services. Unbundling them is another way to build a contrast strategy. Various segments of the market may prefer to get differing subsets of the bundle at better prices. Challengers’ unbundling of the status quo has been facilitated by the internet in one industry after another: Music, TV, and education are leading examples. Incumbents have to make major internal changes to compete with unbundlers, rendering this approach especially effective.

How to begin

1. Precisely identify the assumptions that underlie conventional thinking in your company or industry.
2. Think about what might be gained by proving one or more of them false.
3. Deliberately disturb an aspect of your normal work pattern to break up ingrained assumptions.

What to watch out for

Because the assumptions underlying your business model are embedded in all your processes—and because stable businesses need predictability—it won't be easy to change course. Organizations are very good at resisting change.

Combination: How Can You Connect Products or Services That Have Traditionally Been Separate?

Combination is a canonical creative approach in both the arts and the sciences. As Anthony Brandt and David Eagleman note in *The Runaway Species*, it was by combining two very different ideas—a ride in an elevator and a journey into space—that Albert Einstein found his way to the theory of general relativity. In business, too, creative and successful moves can result from combining things that have been separate. Often these opportunities arise with complementary products and services. Products and payment systems, for example, have traditionally been separate nodes in value chains. But the Chinese social media platform WeChat (owned by Tencent) now includes an integrated mobile payment platform called

WeChat Pay that enables users to buy and sell products within their social networks. Expanding beyond the Chinese ecosystem, Tencent and Alibaba are coordinating with overseas payment firms to enable retailers in other countries to accept their mobile payment services.

Sometimes competitors can benefit from joining forces to grow the pie. (Barry Nalebuff and I explored this idea in our 1996 book *Co-opetition.*) For example, BMW and Daimler have announced plans to combine their mobility services—car sharing, ride hailing, car parking, electric vehicle charging, and tickets for public transport. Presumably, the two automakers hope that this move will be an effective counterattack against Uber and other players that are encroaching on the traditional car industry.

In other instances, companies from wholly separate industries have created new value for customers by combining offerings. Apple and Nike have done so since the 2006 introduction of the Nike+ iPod Sport Kit, which enabled Nike shoes to communicate with an iPod for tracking steps. More recently, versions of the Apple Watch have come with the Nike+ Run Club app fully integrated. Nest Labs and Amazon also complement each other: Nest's intelligent home thermostat becomes even more valuable when it can deploy voice control via Amazon's virtual assistant, Alexa.

New technologies are a rich source of combinatorial possibilities. AI and blockchain come together naturally to

protect the privacy of the large amounts of personal data needed to train algorithms in health care and other sensitive areas. Blockchain and the internet of things come together in the form of sensors and secure data in decentralized applications such as food supply chains, transportation systems, and smart homes, with automated insurance included in smart contracts.

Perhaps the biggest combination today is the one emerging between humans and machines. Some commentators see the future of that relationship as more competitive than cooperative, with humans losing out in many areas of economic life. Others predict a more positive picture, in which machines take on lower-level cognition, freeing humans to be more creative. Martin Reeves and Daichi Ueda have written about algorithms that allow companies to make frequent, calibrated adjustments to their business models, enabling humans to work on high-level objectives and think beyond the present. (See "Designing the Machines That Will Design Strategy," HBR.org, April 2016.)

Strategy from combination involves looking for connections across traditional boundaries, whether by linking a product and a service, two technologies, the upstream and the downstream, or other ingredients. Here, too, the creative strategist must challenge the status quo—this time by thinking not just outside the box but across two or more boxes.

How to begin

1. Form groups with diverse expertise and experience; brainstorm new combinations of products and services.
2. Look for ways to coordinate with providers of complementary products (who may even be competitors).

What to watch out for

Businesses often manage for and measure profits at the individual product or activity level. But combinations require system-level thinking and measurements.

Constraint: How Can You Turn Limitations or Liabilities into Opportunities?

The world's first science fiction story, *Frankenstein*, was written when its author, Mary Wollstonecraft Shelley, was staying near Lake Geneva during an unusually cold and stormy summer and found herself trapped indoors with nothing to do but exercise her imagination. Artists know a lot about constraints—from profound ones, such as serious setbacks in their lives, to structural ones, such as writing a 14-line poem with a specified rhyming structure. In business, too, creative thinking turns limitations into opportunities.

That constraints can spark creative strategies may seem paradoxical. Lift a constraint, and any action that was previously possible is surely still possible; most likely, more is

now possible. But that misses the point that one can think multiple ways in a given situation—and a constraint may prompt a whole new line of thinking. Of course, the Goldilocks principle applies: Too many constraints will choke off all possibilities, and a complete absence of constraints is a problem too.

Tesla hasn't lacked financial resources in entering the car industry, but it doesn't have a traditional dealership network (considered a key part of automakers' business models) through which to sell. Rather than get into the business of building one, Tesla has chosen to sell cars online and to build Apple-like stores staffed with salespeople on salary. This actually positions the company well relative to competitors, whose dealers may be conflicted about promoting electric vehicles over internal-combustion ones. In addition, Tesla controls its pricing directly, whereas consumers who buy electric vehicles from traditional dealers may encounter significant variations in price.

I should note that this attitude toward constraints is very different from that suggested by the classic SWOT analysis. Strategists are supposed to identify the strengths, weaknesses, opportunities, and threats impinging on an organization and then figure out ways to exploit strengths and opportunities and mitigate weaknesses and threats.

In stark contrast, a constraint-based search would look at how those weaknesses could be turned to the company's

advantage. Constraint plus imagination may yield an opportunity.

This approach to strategy turns the SWOT tool upside down in another way as well. Just as an apparent weakness can be turned into a strength, an apparent strength can prove to be a weakness. The likelihood of this often increases over time, as the assets that originally enabled a business to succeed become liabilities when the environment changes. For example, big retailers have historically considered "success" to be moving product out the door; to that end, they needed large physical footprints with on-site inventory. Among the many changes they face today is the rise of "guideshops"—a term used by the menswear retailer Bonobos—where shoppers try on items, which they can have shipped to them or later order online. In the new environment, traditional retail footprints become more of a liability than an asset.

Another way to approach strategy from constraint is to ask whether you might benefit from self-imposed constraints. (Artists do something similar when they choose to work only within a particular medium.) The famous Copenhagen restaurant Noma adheres to the New Nordic Food manifesto (emphasizing purity, simplicity, beauty, seasonality, local tradition, and innovation). A similar strategy of working only with local suppliers has been adopted by thousands of restaurants around the world. A commitment to high environmental standards, fair labor practices, and ethical

supply-chain management can be powerful for organizations looking to lead change in their industries or sectors.

Self-imposed constraints can also spur innovation. Adam Morgan and Mark Barden, in their book *A Beautiful Constraint*, describe the efforts of the Audi racing team in the early 2000s to win Le Mans under the assumption that its cars couldn't go faster than the competition's. Audi developed diesel-powered racers, which required fewer fuel stops than gasoline-powered cars, and won Le Mans three years in succession (2004–2006). In 2017 Audi set itself a new constraint—and a new ambition: to build winning all-electric racers for the new Formula E championship.

How to begin

1. List the "incompetencies" (rather than the competencies) of your organization—and test whether they can in fact be turned into strengths.
2. Consider deliberately imposing some constraints to encourage people to find new ways of thinking and acting.

What to watch out for

Successful businesses face few obvious constraints; people may feel no need to explore how new ones might create new opportunities.

Context: How Can Far-Flung Industries, Ideas, or Disciplines Shed Light on Your Most Pressing Problems?

An entire field, biomimetics, is devoted to finding solutions in nature to problems that arise in engineering, materials science, medicine, and elsewhere. For example, the burrs from the burdock plant, which propagate by attaching to the fur of animals via tiny hooks, inspired George de Mestral in the 1940s to create a clothing fastener that does not jam (as zippers are prone to do). Thus the invention of Velcro. This is a classic problem-solving technique. Start with a problem in one context, find another context in which an analogous problem has already been solved, and import the solution.

Intel did that when it came up with its famous Intel Inside logo, in the early 1990s. The goal was to turn Intel microprocessors into a branded product to speed up consumers' adoption of next-generation chips and, more broadly, to improve the company's ability to drive the PC industry forward. Branded ingredients were well established in certain consumer product sectors—examples include Teflon and NutraSweet—but hadn't been tried in the world of technology. Intel imported the approach to high tech with a novel advertising campaign, successfully branding what had previously been an invisible computer component.

Context switching can be done across industries, as in Intel's case, or even across time. The development of the graphical user interface (GUI) for computers was in a sense the result of a step backward: The developers moved from immersion in the text-based context in which programming had grown up to thinking about the highly visual hand-eye environment in which young children operate. Similarly, some AI researchers are currently looking at how children learn in order to inform processes for machine learning.

Companies are always eager to see into the future, of course, and techniques for trying to do so are well established. That is the purpose of lead-user and extreme-user innovation strategies, which ask companies to shift their attention from mainstream customers to people who are designing their own versions or using products in unexpected ways in especially demanding environments. Information about where the edges of the market are today can signal where the mainstream will be tomorrow. Extreme sports, such as mountain biking, skateboarding, snowboarding, and windsurfing, are good examples. In an MIT Sloan School working paper, Sonali Shah relates that aficionados led many of the innovations in those areas, starting in the 1950s, and big manufacturers added cost efficiencies and marketing to take them mainstream.

When companies locate R&D functions far from headquarters, they're acknowledging the importance of jumping into someone else's context. This is not just a strategy

for large companies that move people to Silicon Valley for tech or the Boston area for biotech. Start-ups, too, should put themselves in the best context for learning and growth. The hardware accelerator HAX, located in Shenzhen, hosts hardware start-up teams from numerous countries and enables them to tap into the high-speed ecosystem of the "hardware capital of the world," quadrupling the rate at which they cycle through iterations of their prototypes.

Strategy focused on context may involve transferring a solution from one setting to another more or less as is. It may mean uncovering entirely new thinking about problems (or opportunities) by finding pioneers who are ahead of the game. At bottom, it's about not being trapped in a single narrative.

How to begin

1. Explain your business to an outsider in another industry. Fresh eyes from a different context can help uncover new answers and opportunities.
2. Engage with lead users, extreme users, and innovation hotspots.

What to watch out for

Businesses need to focus on internal processes to deliver on their current value propositions—but the pressure to focus internally can get in the way of learning from the different contexts in which other players operate.

In the world of management consulting, aspects of "strategy" and "innovation" have started to converge. IDEO, the design and innovation powerhouse, has moved into strategy consulting, for example—while McKinsey has added design-thinking methods to its strategy consulting. This convergence raises an obvious question: If the distinction between strategy and innovation is less clear than it once was, do we really need to think carefully about the role of creativity in the strategy-making process?

I believe strongly that the answer is yes. At its core, strategy is still about finding ways to create and claim value through differentiation. That's a complicated, difficult job. To be sure, it requires tools that can help identify surprising, creative breaks from conventional thinking. But it also requires tools for analyzing the competitive landscape, the dynamics threatening that landscape, and a company's resources and competencies. We need to teach business school students—and executives—how to be creative and rigorous at the same time.

Originally published in March–April 2019. Reprint R1902C

What Most People Get Wrong about Men and Women

by Catherine H. Tinsley and Robin J. Ely

THE CONVERSATION ABOUT the treatment of women in the workplace has reached a crescendo of late, and senior leaders—men as well as women—are increasingly vocal about a commitment to gender parity. That's all well and good, but there's an important catch. The discussions, and many of the initiatives companies have undertaken, too often reflect a faulty belief: that men and women are fundamentally *different*, by virtue of their genes or their upbringing or both. Of course, there are biological differences. But those are not the differences people are usually talking about. Instead, the rhetoric focuses on the idea that women are inherently unlike men in terms of disposition, attitudes, and behaviors. (Think headlines that tout "Why women do X at the office" or "Working women don't Y.")

One set of assumed differences is marshaled to explain women's failure to achieve parity with men: Women negotiate

poorly, lack confidence, are too risk-averse, or don't put in the requisite hours at work because they value family more than their careers. Simultaneously, other assumed differences—that women are more caring, cooperative, or mission-driven—are used as a rationale for companies to invest in women's success. But whether framed as a barrier or a benefit, these beliefs hold women back. We will not level the playing field so long as the bedrock on which it rests is our conviction about how the sexes are different.

The reason is simple: Science, by and large, does not actually support these claims. There is wide variation among women and among men, and meta-analyses show that, on average, the sexes are far more similar in their inclinations, attitudes, and skills than popular opinion would have us believe. We do see sex differences in various settings, including the workplace—but those differences are not rooted in fixed gender traits. Rather, they stem from organizational structures, company practices, and patterns of interaction that position men and women differently, creating systematically different experiences for them. When facing dissimilar circumstances, people respond differently—not because of their sex but because of their situations.

Emphasizing sex differences runs the risk of making them seem natural and inevitable. As anecdotes that align with stereotypes are told and retold, without addressing why and when stereotypical behaviors appear, sex differences are

exaggerated and take on a determinative quality. Well-meaning but largely ineffectual interventions then focus on "fixing" women or accommodating them rather than on changing the circumstances that gave rise to different behaviors in the first place.

Take, for example, the common belief that women are more committed to family than men are. Research simply does not support that notion. In a study of Harvard Business School graduates that one of us conducted, nearly everyone, regardless of gender, placed a higher value on their families than on their work (see "Rethink What You 'Know' About High-Achieving Women," HBR, December 2014). Moreover, having made career decisions to accommodate family responsibilities didn't explain the gender achievement gap. Other research, too, makes it clear that men and women do not have fundamentally different priorities.

Idea in Brief

The Belief

There's a popular notion that men and women are fundamentally different in important (nonbiological) ways—and those differences are cited to explain women's lagged achievement.

The Truth

According to numerous meta-analyses of published research, men and women are actually very similar with respect to key attributes such as confidence, appetite for risk, and negotiating skill.

Why It Matters

> Too many managers try to “fix” women or accommodate their supposed differences—and that doesn’t work. Companies must instead address the organizational conditions that lead to lower rates of retention and promotion for women.

Numerous studies show that what does differ is the treatment mothers and fathers receive when they start a family. Women (but not men) are seen as needing support, whereas men are more likely to get the message—either explicit or subtle—that they need to “man up” and not voice stress and fatigue. If men do ask, say, for a lighter travel schedule, their supervisors may cut them some slack—but often grudgingly and with the clear expectation that the reprieve is temporary. Accordingly, some men attempt an under-the-radar approach, quietly reducing hours or travel and hoping it goes unnoticed, while others simply concede, limiting the time they spend on family responsibilities and doubling down at work. Either way, they maintain a reputation that keeps them on an upward trajectory. Meanwhile, mothers are often expected, indeed encouraged, to ratchet back at work. They are rerouted into less taxing roles and given less “demanding” (read: lower-status, less career-enhancing) clients.

To sum up, men’s and women’s desires and challenges about work/family balance are remarkably similar. It is what they experience at work once they become parents that puts them in very different places.

Things don't have to be this way. When companies observe differences in the overall success rates of women and men, or in behaviors that are critical to effectiveness, they can actively seek to understand the organizational conditions that might be responsible, and then they can experiment with changing those conditions.

Consider the example of a savvy managing director concerned about the leaky pipeline at her professional services firm. Skeptical that women were simply "opting out" following the birth of a child, she investigated and found that one reason women were leaving the firm stemmed from the performance appraisal system: Supervisors had to adhere to a forced distribution when rating their direct reports, and women who had taken parental leave were unlikely to receive the highest rating because their performance was ranked against that of peers who had worked a full year. Getting less than top marks not only hurt their chances of promotion but also sent a demoralizing message that being a mother was incompatible with being on a partner track. However, the fix was relatively easy: The company decided to reserve the forced distribution for employees who worked the full year, while those with long leaves could roll over their rating from the prior year. That applied to both men and women, but the policy was most heavily used by new mothers. The change gave women more incentive to return from maternity leave and helped keep them on track for advancement. Having more mothers stay on track,

in turn, helped chip away at assumptions within the firm about women's work/family preferences.

As this example reveals, companies need to dive deeper into their beliefs, norms, practices, and policies to understand how they position women relative to men and how the different positions fuel inequality. Seriously investigating the context that gives rise to differential patterns in the way men and women experience the workplace—and intervening accordingly—can help companies chart a path to gender parity.

Below, we address three popular myths about how the sexes differ and explain how each manifests itself in organizational discourse about women's lagged advancement. Drawing on years of social science research, we debunk the myths and offer alternative explanations for observed sex differences—explanations that point to ways that managers can level the playing field. We then offer a four-pronged strategy for undertaking such actions.

Popular Myths

We've all heard statements in the media and in companies that women lack *the desire or ability to negotiate*, that they lack *confidence*, and that they lack *an appetite for risk*. And, the thinking goes, those shortcomings explain why women have so far failed to reach parity with men.

For decades, studies have examined sex differences on these three dimensions, enabling social scientists to conduct meta-analyses—investigations that reveal whether or not, on average across studies, sex differences hold, and if so, how large the differences are. (See the sidebar "The Power of Meta-Analysis.") Just as importantly, meta-analyses also reveal the circumstances under which differences between men and women are more or less likely to arise. The aggregated findings are clear: Context explains any sex differences that exist in the workplace.

Take negotiation. Over and over, we hear that women are poor negotiators—they "settle too easily," are "too nice," or are "too cooperative." But not so, according to research. Jens Mazei and colleagues recently analyzed more than 100 studies examining whether men and women negotiate different outcomes; they determined that gender differences were small to negligible. Men have a slight advantage in negotiations when they are advocating exclusively for themselves and when ambiguity about the stakes or opportunities is high. Larger disparities in outcomes occur when negotiators either have no prior experience or are forced to negotiate, as in a mandated training exercise. But such situations are atypical, and even when they do arise, statisticians would deem the resulting sex differences to be small. As for the notion that women are more cooperative than men, research by Daniel Balliet and colleagues refutes that.

The belief that women lack confidence is another fallacy. That assertion is commonly invoked to explain why women speak up less in meetings and do not put themselves forward for promotions unless they are 100% certain they meet all the job requirements. But research does not corroborate the idea that women are less confident than men. Analyzing more than 200 studies, Kristen Kling and colleagues concluded that the only noticeable differences occurred during adolescence; starting at age 23, differences become negligible.

What about risk taking—are women really more conservative than men? Many people believe that's true—though they are split on whether being risk-averse is a strength or a weakness. On the positive side, the thinking goes, women are less likely to get caught up in macho displays of bluff and bravado and thus are less likely to take unnecessary risks. Consider the oft-heard sentiment following the demise of Lehman Brothers: "If Lehman Brothers had been Lehman Sisters, the financial crisis might have been averted." On the negative side, women are judged as too cautious to make high-risk, potentially high-payoff investments.

But once again, research fails to support either of these stereotypes. As with negotiation, sex differences in the propensity to take risks are small and depend on the context. In a meta-analysis performed by James Byrnes and colleagues, the largest differences arise in contexts unlikely to exist in most organizations (such as among people asked to participate

in a game of pure chance). Similarly, in a study Peggy Dwyer and colleagues ran examining the largest, last, and riskiest investments made by nearly 2,000 mutual fund investors, sex differences were very small. More importantly, when investors' specific knowledge about the investments was added to the equation, the sex difference diminished to near extinction, suggesting that access to information, not propensity for risk taking, explains the small sex differences that have been documented.

In short, a wealth of evidence contradicts each of these popular myths. Yet they live on through oft-repeated narratives routinely invoked to explain women's lagged advancement.

More-Plausible Explanations

The extent to which employees are able to thrive and succeed at work depends partly on the kinds of opportunities and treatment they receive. People are more likely to behave in ways that undermine their chances for success when they are disconnected from information networks, when they are judged or penalized disproportionately harshly for mistakes or failures, and when they lack feedback. Unfortunately, women are more likely than men to encounter each of these situations. And the way they respond—whether that's by failing to drive a hard bargain, to speak up, or to take risks—gets unfairly

attributed to “the way women are,” when in fact the culprit is very likely the differential conditions they face.

Why the Sex-Difference Narrative Persists

BELIEFS IN SEX DIFFERENCES have staying power partly because they uphold conventional gender norms, preserve the gender status quo, and require no upheaval of existing organizational practices or work arrangements. But they are also the path of least resistance for our brains. Three well-documented cognitive errors help explain the endurance of the sex-difference narrative.

First, when seeking to explain others’ behavior, we gravitate to explanations based on intrinsic *personality traits*—including stereotypically “male” traits and stereotypically “female” traits”—rather than *contextual factors*. (Social psychologists call this “the fundamental attribution error.”) For example, if a man speaks often and forcefully in a meeting, we are more likely to conclude that he is assertive and confident than to search for a situational explanation, such as that he’s been repeatedly praised for his contributions. Likewise, if a woman is quiet in a meeting, the easier explanation is that she’s meek or underconfident; it takes more cognitive energy to construct an alternative account, such as that she is used to being cut off or ignored when she speaks. In short, when we see men and women behaving in gender-stereotypical ways, we tend to make the most cognitively simple assumption—that the behavior reflects who they are rather than the situation they are in.

Second, mere exposure to a continuing refrain, such as “Women are X, and men are Y,” makes people judge the statement as true. Many beliefs—that bats are blind, that fresh produce is always more nutritious than frozen, that you shouldn’t wake a sleepwalker—are

repeated so often that their mere familiarity makes them easier for our minds to accept as truth. (This is called the "mere exposure effect.")

Third, once people believe something is true, they tend to seek, notice, and remember evidence that confirms the position and to ignore or forget evidence that would challenge it. (Psychologists call this "confirmation bias.") If we believe that gender stereotypes are accurate, we are more likely to expect, notice, and remember times when men and women behave in gender-stereotypical ways and to overlook times when they don't.

Multiple studies show, for example, that women are less embedded in networks that offer opportunities to gather vital information and garner support. When people lack access to useful contacts and information, they face a disadvantage in negotiations. They may not know what is on the table, what is within the realm of possibility, or even that a chance to strike a deal exists. When operating under such conditions, women are more likely to conform to the gender stereotype that "women don't ask."

We saw this dynamic vividly play out when comparing the experiences of two professionals we'll call Mary and Rick. (In this example and others that follow, we have changed the names and some details to maintain confidentiality.) Mary and Rick were both midlevel advisers in the wealth management division of a financial services firm. Rick was able to bring in more assets to manage because he sat on the board of a nonprofit, giving him access to a pool of potential clients with

high net worth. What Mary did not know for many years is how Rick had gained that advantage. Through casual conversations with one of the firm's senior partners, with whom he regularly played tennis, Rick had learned that discretionary funds existed to help advisers cultivate relationships with clients. So he arranged for the firm to make a donation to the nonprofit. He then began attending the nonprofit's fund-raising events and hobnobbing with key players, eventually parlaying his connections into a seat on the board. Mary, by contrast, had no informal relationships with senior partners at the firm and no knowledge of the level of resources that could have helped her land clients.

When people are less embedded, they are also less aware of opportunities for stretch assignments and promotions, and their supervisors may be in the dark about their ambitions. But when women fail to "lean in" and seek growth opportunities, it is easy to assume that they lack the confidence to do so—not that they lack pertinent information. Julie's experience is illustrative. Currently the CEO of a major investment fund, Julie had left her previous employer of 15 years after learning that a more junior male colleague had leapfrogged over her to fill an opening she didn't even know existed. When she announced that she was leaving and why, her boss was surprised. He told her that if he had realized she wanted to move up, he would have gladly helped position her for the promotion. But because she hadn't put her hat in the ring, he

had assumed she lacked confidence in her ability to handle the job.

How people react to someone's mistake or failure can also affect that person's ability to thrive and succeed. Several studies have found that because women operate under a higher-resolution microscope than their male counterparts do, their mistakes and failures are scrutinized more carefully and punished more severely. People who are scrutinized more carefully will, in turn, be less likely to speak up in meetings, particularly if they feel no one has their back. However, when women fail to speak up, it is commonly assumed that they lack confidence in their ideas.

We saw a classic example of this dynamic at a biotech company in which team leaders noticed that their female colleagues, all highly qualified research scientists, participated far less in team meetings than their male counterparts did, yet later, in one-on-one conversations, often offered insightful ideas germane to the discussion. What these leaders had failed to see was that when women did speak in meetings, their ideas tended to be either ignored until a man restated them or shot down quickly if they contained even the slightest flaw. In contrast, when men's ideas were flawed, the meritorious elements were salvaged. Women therefore felt they needed to be 110% sure of their ideas before they would venture to share them. In a context in which being smart was the coin of the

realm, it seemed better to remain silent than to have one's ideas repeatedly dismissed.

It stands to reason that people whose missteps are more likely to be held against them will also be less likely to take risks. That was the case at a Big Four accounting firm that asked us to investigate why so few women partners were in formal leadership roles. The reason, many believed, was that women did not want such roles because of their family responsibilities, but our survey revealed a more complex story. First, women and men were equally likely to say they would accept a leadership role if offered one, but men were nearly 50% more likely to have been offered one. Second, women were more likely than men to say that worries about jeopardizing their careers deterred them from pursuing leadership positions—they feared they would not recover from failure and thus could not afford to take the risks an effective leader would need to take. Research confirms that such concerns are valid. For example, studies by Victoria Brescoll and colleagues found that if women in male-dominated occupations make mistakes, they are accorded less status and seen as less competent than men making the same mistakes; a study by Ashleigh Rosette and Robert Livingston demonstrated that black women leaders are especially vulnerable to this bias.

Research also shows that women get less frequent and lower-quality feedback than men. When people don't receive feedback, they are less likely to know their worth in

negotiations. Moreover, people who receive little feedback are ill-equipped to assess their strengths, shore up their weaknesses, and judge their prospects for success and are therefore less able to build the confidence they need to proactively seek promotions or make risky decisions.

An example of this dynamic comes from a consulting firm in which HR staff members delivered partners' annual feedback to associates. The HR folks noticed that when women were told they were "doing fine," they "freaked out," feeling damned by faint praise; when men received the same feedback, they left the meeting "feeling great." HR concluded that women lack self-confidence and are therefore more sensitive to feedback, so the team advised partners to be especially encouraging to the women associates and to soften any criticism. Many of the partners were none too pleased to have to treat a subset of their associates with kid gloves, grousing that "if women can't stand the heat, they should get out of the kitchen." What these partners failed to realize, however, is that the kitchen was a lot hotter for women in the firm than for men. Why? Because the partners felt more comfortable with the men and so were systematically giving them more informal, day-to-day feedback. When women heard in their annual review that they were doing "fine," it was often the first feedback they'd received all year; they had nothing else to go on and assumed it meant their performance was merely adequate. In contrast, when men heard they were doing "fine," it was but one piece of

information amidst a steady stream. The upshot was disproportionate turnover among women associates, many of whom left the firm because they believed their prospects for promotion were slim.

An Alternative Approach

The problem with the sex-difference narrative is that it leads companies to put resources into "fixing" women, which means that women miss out on what they need—and what every employee deserves: a context that enables them to reach their potential and maximizes their chances to succeed.

Managers who are advancing gender equity in their firms are taking a more inquisitive approach—rejecting old scripts, seeking an evidence-based understanding of how women experience the workplace, and then creating the conditions that increase women's prospects for success. Their approach entails four steps:

1. Question the narrative

A consulting firm we worked with had recruited significant numbers of talented women into its entry ranks—and then struggled to promote them. Their supervisors' explanations? Women are insufficiently competitive, lack "fire in the belly," or don't have the requisite confidence to excel in the job. But those narratives did not ring true to Sarah, a regional head, because a handful of women—those within her region—were

performing and advancing at par. So rather than accept her colleagues' explanations, she got curious.

2. Generate a plausible alternative explanation

Sarah investigated the factors that might have helped women in her region succeed and found that they received more hands-on training and more attention from supervisors than did women in other regions. This finding suggested that the problem lay not with women's deficiencies but with their differential access to the conditions that enhance self-confidence and success.

To test that hypothesis, Sarah designed an experiment, with our help. First, we randomly split 60 supervisors into two groups of 30 for a training session on coaching junior consultants. Trainers gave both groups the same lecture on how to be a good coach. With one group, however, trainers shared research showing that differences in men's and women's self-confidence are minuscule, thus subtly giving the members of this "treatment" group reason to question gender stereotypes. The "control" group didn't get that information. Next, trainers gave all participants a series of hypotheticals in which an employee—sometimes a man and sometimes a woman—was underperforming. In both groups, participants were asked to write down the feedback they would give the underperforming employee.

Clear differences emerged between the two groups. Supervisors in the control group took different tacks with the underperforming man and woman: They were far less critical of the woman and focused largely on making her feel good, whereas they gave the man feedback that was more direct, specific, and critical, often with concrete suggestions for how he could improve. In contrast, the supervisors who had been shown research that refuted sex differences in self-confidence gave both employees the same kind of feedback; they also asked for more-granular information about the employee's performance so that they could deliver constructive comments. We were struck by how the participants who had been given a reason to question gender stereotypes focused on learning more about individuals' specific performance problems.

The experiment confirmed Sarah's sense that women's lagged advancement might be due at least partly to supervisors' assumptions about the training and development needs of their female direct reports. Moreover, her findings gave supervisors a plausible alternative explanation for women's lagged advancement—a necessary precondition for taking the next step. Although different firms find different types of evidence more or less compelling—not all require as rigorous a test as this firm did—Sarah's evidence-based approach illustrates a key part of the strategy we are advocating.

3. Change the context and assess the results

Once a plausible alternative explanation has been developed, companies can make appropriate changes and see if performance improves. Two stories help illustrate this step. Both come from a midmarket private equity firm that was trying to address a problem that had persisted for 10 years: The company's promotion and retention rates for white women and people of color were far lower than its hiring rates.

The first story involves Elaine, an Asian-American senior associate who wanted to sharpen her financing skills and asked Dave, a partner, if she could assist with that aspect of his next deal. He invited her to lunch, but when they met, he was underwhelmed. Elaine struck him as insufficiently assertive and overly cautious. He decided against putting her on his team—but then he had second thoughts. The partners had been questioning their ability to spot and develop talent, especially in the case of associates who didn't look like them. Dave thus decided to try an experiment: He invited Elaine to join the team and then made a conscious effort to treat her exactly as he would have treated someone he deemed a superstar. He introduced her to the relevant players in the industry, told the banks she would be leading the financing, and gave her lots of rope but also enough feedback and coaching so that she wouldn't hang herself. Elaine did not disappoint; indeed, her performance was stellar. While quiet in demeanor, Dave's new protégée showed an uncanny ability to

read the client and come up with creative approaches to the deal's financing.

A second example involves Ned, a partner who was frustrated that Joan, a recent-MBA hire on his team, didn't assert herself on management team calls. At first Ned simply assumed that Joan lacked confidence. But then it occurred to him that he might be falling back on gender stereotypes, and he took a closer look at his own behavior. He realized that he wasn't doing anything to make participation easier for her and was actually doing things that made it harder, like taking up all the airtime on calls. So they talked about it, and Joan admitted that she was afraid of making a mistake and was hyperaware that if she spoke, she needed to say something very smart. Ned realized that he, too, was afraid she would make a mistake or wouldn't add value to the discussion, which is partly why he took over. But on reflection, he saw that it wouldn't be the end of the world if she did stumble—he did the same himself now and again. For their next few calls, they went over the agenda beforehand and worked out which parts she would take the lead on; he then gave her feedback after the call. Ned now has a junior colleague to whom he can delegate more; Joan, meanwhile, feels more confident and has learned that she can take risks and recover from mistakes.

The Power of Meta-Analysis

A META-ANALYSIS is a statistical technique used to combine the results of many studies, providing a more reliable basis for drawing conclusions from research. This approach has three advantages over a single study.

First, it is more *accurate*, because it is based on a very large sample—the total of the samples across all the studies—and because it contains data collected in many different contexts. Any single set of findings may reflect idiosyncrasies of the study's sample or context and thus may not yield conclusions that are truly generalizable. A meta-analysis, in essence, averages across these idiosyncrasies to give us a truer answer to the research question (in this case, "Are men and women different with regard to a particular trait or behavior?").

Second, a meta-analysis is more *comprehensive*. Because it contains studies conducted in many different contexts, it can tell us in which kinds of contexts we are more or less likely to see sex differences.

Third, a meta-analysis is more *precise:* It can tell us just how different men and women are. For any given trait or behavior, there is variability *among* men and *among* women; typically, those within-group differences are distributed around some "true" average for each group. Using the averages and the variability within each group, we can calculate an "effect size" that can be thought of as the impact that sex has on a particular trait. When testing for a sex difference, we are in essence asking the question "How much overlap is there between women and men, or, stated another way, how far apart are their respective averages, relative to the variability within each sex?"

Take the left-hand graph, which shows the distribution of men's and women's heights in the UK. We can see from the curves that men, on average, are quite a bit taller than women. In fact, men average five feet, nine inches, and women five feet, three inches—a six-inch difference. We can also see that a number of women are taller than the average man, just as a number of men are shorter than the average woman. The size of the sex effect on height is 1.72, which is considered "large."

Using that sex difference as a reference point, we can see from the right-hand graph that the difference between men and women in self-esteem, or confidence, is much smaller, with an effect size of 0.10. Although the difference in each graph is statistically significant, the difference in confidence is considered, from a statistical point of view, "trivial"—and from a managerial point of view, essentially meaningless. This same analysis for men's and women's negotiation outcomes and for their propensity to take risks yielded effect sizes of 0.20 ("small") and 0.13 ("trivial"), respectively. In short, contrary to popular belief, all three sex differences we consider in this article are, for all intents and purposes, meaningless.

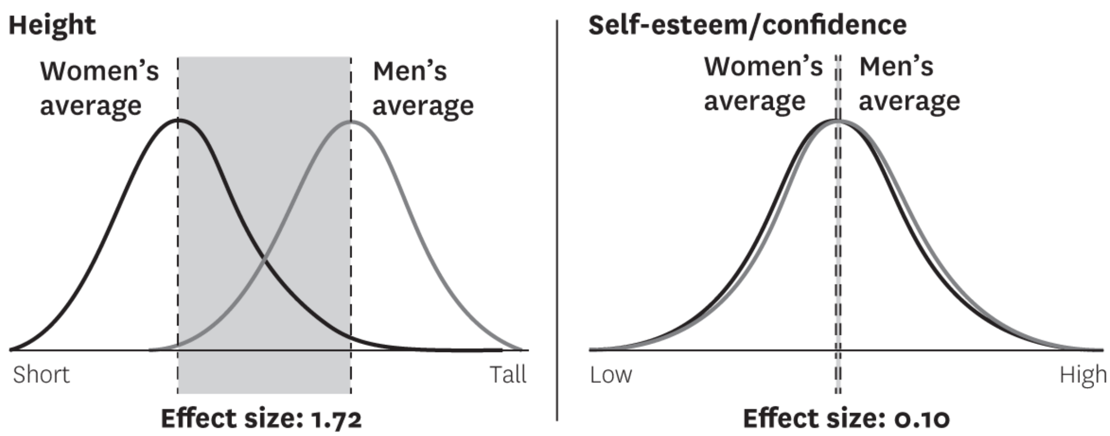

Note: Statisticians consider an effect size of less than 0.20 to be "trivial," 0.20–0.49 to be "small," 0.50–0.79 to be "medium," and 0.80 or more to be "large."

4. Promote continual learning

Both Dave and Ned recognized that their tendency to jump to conclusions based on stereotypes was robbing them—and the firm—of vital talent. Moreover, they have seen firsthand how questioning assumptions and proactively changing conditions gives women the opportunity to develop and shine. The lessons from these small-scale experiments are ongoing: Partners at the firm now meet regularly to discuss what they're

learning. They also hold one another accountable for questioning and testing gender-stereotypical assessments as they arise. As a result, old narratives about women's limitations are beginning to give way to new narratives about how the firm can better support all employees.

The four steps we've outlined are consistent with research suggesting that on difficult issues such as gender and race, managers respond more positively when they see themselves as part of the solution rather than simply part of the problem. The solution to women's lagged advancement is not to fix women or their managers but to fix the conditions that undermine women and reinforce gender stereotypes. Furthermore, by taking an inquisitive, evidence-based approach to understanding behavior, companies can not only address gender disparities but also cultivate a learning orientation and a culture that gives all employees the opportunity to reach their full potential.

Originally published in May–June 2018. Reprint R1803J

Collaborative Intelligence

Humans and AI Are Joining Forces.

by H. James Wilson and Paul R. Daugherty

ARTIFICIAL INTELLIGENCE IS BECOMING good at many "human" jobs—diagnosing disease, translating languages, providing customer service—and it's improving fast. This is raising reasonable fears that AI will ultimately replace human workers throughout the economy. But that's not the inevitable, or even most likely, outcome. Never before have digital tools been so responsive to us, nor we to our tools. While AI will radically alter how work gets done and who does it, the technology's larger impact will be in complementing and augmenting human capabilities, not replacing them.

Certainly, many companies have used AI to automate processes, but those that deploy it mainly to displace employees will see only short-term productivity gains. In our research involving 1,500 companies, we found that firms achieve the most significant performance improvements when humans and machines work together (see the exhibit "The value of collaboration"). Through such collaborative intelligence, humans

and AI actively enhance each other's complementary strengths: the leadership, teamwork, creativity, and social skills of the former, and the speed, scalability, and quantitative capabilities of the latter. What comes naturally to people (making a joke, for example) can be tricky for machines, and what's straightforward for machines (analyzing gigabytes of data) remains virtually impossible for humans. Business requires both kinds of capabilities.

The value of collaboration

Companies benefit from optimizing collaboration between humans and artificial intelligence. Five principles can help them do so: Reimagine business processes; embrace experimentation/employee involvement; actively direct AI strategy; responsibly collect data; and redesign work to incorporate AI and cultivate related employee skills. A survey of 1,075 companies in 12 industries found that the more of these principles companies adopted, the better their AI initiatives performed in terms of speed, cost savings, revenues, or other operational measures.

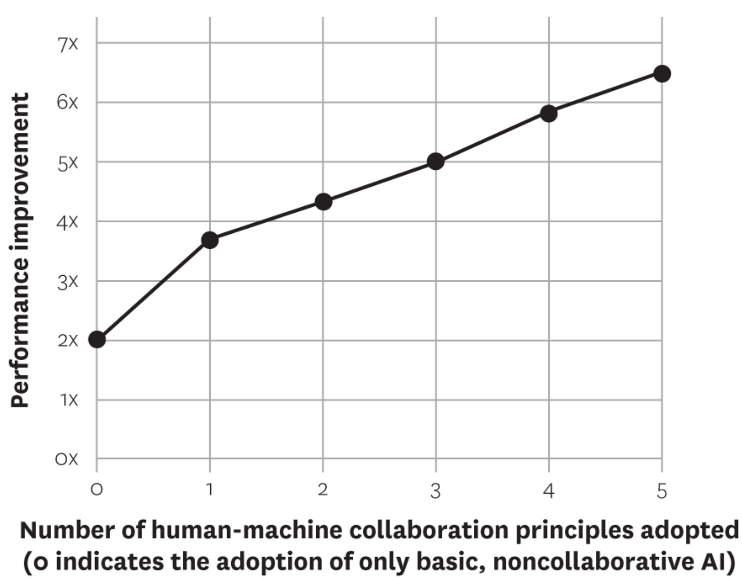

To take full advantage of this collaboration, companies must understand how humans can most effectively augment machines, how machines can enhance what humans do best, and how to redesign business processes to support the partnership. Through our research and work in the field, we have developed guidelines to help companies achieve this and put the power of collaborative intelligence to work.

Idea in Brief

The Outlook

Artificial intelligence is transforming business—and having the most significant impact when it augments human workers instead of replacing them.

The Details

Companies see the biggest performance gains when humans and smart machines collaborate. People are needed to train machines, explain their outputs, and ensure their responsible use. AI, in turn, can enhance humans' cognitive skills and creativity, free workers from low-level tasks, and extend their physical capabilities.

The Prescription

Companies should reimagine their business processes, focusing on using AI to achieve more operational flexibility or speed, greater scale, better decision making, or increased personalization of products and services.

Humans Assisting Machines

Humans need to perform three crucial roles. They must *train* machines to perform certain tasks; *explain* the outcomes of those

tasks, especially when the results are counterintuitive or controversial; and *sustain* the responsible use of machines (by, for example, preventing robots from harming humans).

Training

Machine-learning algorithms must be taught how to perform the work they're designed to do. In that effort, huge training data sets are amassed to teach machine-translation apps to handle idiomatic expressions, medical apps to detect disease, and recommendation engines to support financial decision making. In addition, AI systems must be trained how best to interact with humans. While organizations across sectors are now in the early stages of filling trainer roles, leading tech companies and research groups already have mature training staffs and expertise.

Consider Microsoft's AI assistant, Cortana. The bot required extensive training to develop just the right personality: confident, caring, and helpful but not bossy. Instilling those qualities took countless hours of attention by a team that included a poet, a novelist, and a playwright. Similarly, human trainers were needed to develop the personalities of Apple's Siri and Amazon's Alexa to ensure that they accurately reflected their companies' brands. Siri, for example, has just a touch of sassiness, as consumers might expect from Apple.

AI assistants are now being trained to display even more complex and subtle human traits, such as sympathy. The start-up Koko, an offshoot of the MIT Media Lab, has developed

technology that can help AI assistants seem to commiserate. For instance, if a user is having a bad day, the Koko system doesn't reply with a canned response such as "I'm sorry to hear that." Instead it may ask for more information and then offer advice to help the person see his issues in a different light. If he were feeling stressed, for instance, Koko might recommend thinking of that tension as a positive emotion that could be channeled into action.

Explaining

As AIs increasingly reach conclusions through processes that are opaque (the so-called black-box problem), they require human experts in the field to explain their behavior to nonexpert users. These "explainers" are particularly important in evidence-based industries, such as law and medicine, where a practitioner needs to understand how an AI weighed inputs into, say, a sentencing or medical recommendation. Explainers are similarly important in helping insurers and law enforcement understand why an autonomous car took actions that led to an accident—or failed to avoid one. And explainers are becoming integral in regulated industries—indeed, in any consumer-facing industry where a machine's output could be challenged as unfair, illegal, or just plain wrong. For instance, the European Union's new General Data Protection Regulation (GDPR) gives consumers the right to receive an explanation for any algorithm-based decision, such as the rate offer on a credit card or mortgage. This is one area where AI will contribute to *increased* employment: Experts estimate

that companies will have to create about 75,000 new jobs to administer the GDPR requirements.

Sustaining

In addition to having people who can explain AI outcomes, companies need "sustainers"—employees who continually work to ensure that AI systems are functioning properly, safely, and responsibly.

For example, an array of experts sometimes referred to as safety engineers focus on anticipating and trying to prevent harm by AIs. The developers of industrial robots that work alongside people have paid careful attention to ensuring that they recognize humans nearby and don't endanger them. These experts may also review analysis from explainers when AIs do cause harm, as when a self-driving car is involved in a fatal accident.

Other groups of sustainers make sure that AI systems uphold ethical norms. If an AI system for credit approval, for example, is found to be discriminating against people in certain groups (as has happened), these ethics managers are responsible for investigating and addressing the problem. Playing a similar role, data compliance officers try to ensure that the data that is feeding AI systems complies with the GDPR and other consumer-protection regulations. A related data-use role involves ensuring that AIs manage information responsibly. Like many tech companies, Apple uses AI to collect personal details about users as they engage with the company's devices and software. The

aim is to improve the user experience, but unconstrained data gathering can compromise privacy, anger customers, and run afoul of the law. The company's "differential privacy team" works to make sure that while the AI seeks to learn as much as possible about a group of users in a statistical sense, it is protecting the privacy of individual users.

Machines Assisting Humans

Smart machines are helping humans expand their abilities in three ways. They can *amplify* our cognitive strengths; *interact* with customers and employees to free us for higher-level tasks; and *embody* human skills to extend our physical capabilities.

Enhancing performance

At organizations in all kinds of industries, humans and AI are collaborating to improve five elements of business processes.

Flexibility	Auto manufacturing	Mercedes-Benz	Assembly robots work safely alongside humans to customize cars in real time.
	Product design	Autodesk	Software suggests new product design concepts as a designer changes parameters such as materials, cost, and performance requirements.
	Software development	Gigster	AI helps analyze any type of software project, no matter the size or complexity, enabling humans to quickly estimate the work required, organize experts, and adapt workflows in real time.
Speed	Fraud detection	HSBC	AI screens credit- and debit-card transactions to instantly approve legitimate ones while flagging questionable ones for humans to evaluate.
	Cancer treatment	Roche	AI aggregates patient data from disparate IT systems, speeding collaboration among specialists.
	Public safety	Singapore government	Video analytics during public events predicts crowd behavior, helping responders address security incidents rapidly.
Scale	Recruiting	Unilever	Automated applicant screening dramatically expands the pool of qualified candidates for hiring managers to evaluate.
	Customer service	Virgin Trains	Bot responds to basic customer requests, doubling the volume handled and freeing humans to address more-complex issues.
	Casino management	GGH Morowitz	Computer-vision system helps humans continuously monitor every gaming table in a casino.
Decision making	Equipment maintenance	General Electric	"Digital twins" and Predix diagnostic application provide techs with tailored recommendations for machine maintenance.
	Financial services	Morgan Stanley	Robo-advisers offer clients a range of investment options based on real-time market information.
	Disease prediction	Icahn School of Medicine at Mount Sinai	Deep Patient system helps doctors predict patients' risk of specific disease, allowing preventive intervention.
Personalization	Guest experience	Carnival Corporation	Wearable AI device streamlines the logistics of cruise-ship activities and anticipates guest preferences, facilitating tailored staff support.
	Health care	Pfizer	Wearable sensors for Parkinson's patients track symptoms 24/7, allowing customized treatment.
	Retail fashion	Stitch Fix	AI analyzes customer data to advise human stylists, who give customers individualized clothing and styling recommendations.

Amplifying

Artificial intelligence can boost our analytic and decision-making abilities by providing the right information at the right time. But it can also heighten creativity. Consider how Autodesk's

Dreamcatcher AI enhances the imagination of even exceptional designers. A designer provides Dreamcatcher with criteria about the desired product—for example, a chair able to support up to 300 pounds, with a seat 18 inches off the ground, made of materials costing less than $75, and so on. She can also supply information about other chairs that she finds attractive. Dreamcatcher then churns out thousands of designs that match those criteria, often sparking ideas that the designer might not have initially considered. She can then guide the software, telling it which chairs she likes or doesn't, leading to a new round of designs.

Throughout the iterative process, Dreamcatcher performs the myriad calculations needed to ensure that each proposed design meets the specified criteria. This frees the designer to concentrate on deploying uniquely human strengths: professional judgment and aesthetic sensibilities.

Interacting

Human-machine collaboration enables companies to interact with employees and customers in novel, more effective ways. AI agents like Cortana, for example, can facilitate communications between people or on behalf of people, such as by transcribing a meeting and distributing a voice-searchable version to those who couldn't attend. Such applications are inherently scalable—a single chatbot, for instance, can provide routine customer service to large numbers of people simultaneously, wherever they may be.

SEB, a major Swedish bank, now uses a virtual assistant called Aida to interact with millions of customers. Able to handle natural-language conversations, Aida has access to vast stores of data and can answer many frequently asked questions, such as how to open an account or make cross-border payments. She can also ask callers follow-up questions to solve their problems, and she's able to analyze a caller's tone of voice (frustrated versus appreciative, for instance) and use that information to provide better service later. Whenever the system can't resolve an issue—which happens in about 30% of cases—it turns the caller over to a human customer-service representative and then monitors that interaction to learn how to resolve similar problems in the future. With Aida handling basic requests, human reps can concentrate on addressing more-complex issues, especially those from unhappy callers who might require extra hand-holding.

Embodying

Many AIs, like Aida and Cortana, exist principally as digital entities, but in other applications the intelligence is embodied in a robot that augments a human worker. With their sophisticated sensors, motors, and actuators, AI-enabled machines can now recognize people and objects and work safely alongside humans in factories, warehouses, and laboratories.

In manufacturing, for example, robots are evolving from potentially dangerous and "dumb" industrial machines into smart, context-aware "cobots." A cobot arm might, for example, handle repetitive actions that require heavy lifting, while a

person performs complementary tasks that require dexterity and human judgment, such as assembling a gear motor.

Hyundai is extending the cobot concept with exoskeletons. These wearable robotic devices, which adapt to the user and location in real time, will enable industrial workers to perform their jobs with superhuman endurance and strength.

Reimagining Your Business

In order to get the most value from AI, operations need to be redesigned. To do this, companies must first discover and describe an operational area that can be improved. It might be a balky internal process (such as HR's slowness to fill staff positions), or it could be a previously intractable problem that can now be addressed using AI (such as quickly identifying adverse drug reactions across patient populations). Moreover, a number of new AI and advanced analytic techniques can help surface previously invisible problems that are amenable to AI solutions (see the sidebar "Revealing Invisible Problems").

Revealing Invisible Problems

FORMER U.S. DEFENSE SECRETARY Donald Rumsfeld once famously distinguished among "known knowns," "known unknowns," and "unknown unknowns"—things you're not even aware you don't know. Some companies are now using AI to uncover unknown unknowns in their businesses. Case in point: GNS Healthcare applies machine-learning software to find overlooked relationships among data in

patients' health records and elsewhere. After identifying a relationship, the software churns out numerous hypotheses to explain it and then suggests which of those are the most likely. This approach enabled GNS to uncover a new drug interaction hidden in unstructured patient notes. CEO Colin Hill points out that this is not garden-variety data mining to find associations. "Our machine-learning platform is not just about seeing patterns and correlations in data," he says. "It's about actually discovering causal links."

Next, companies must develop a solution through co-creation —having stakeholders envision how they might collaborate with AI systems to improve a process. Consider the case of a large agricultural company that wanted to deploy AI technology to help farmers. An enormous amount of data was available about soil properties, weather patterns, historical harvests, and so forth, and the initial plan was to build an AI application that would more accurately predict future crop yields. But in discussions with farmers, the company learned of a more pressing need. What farmers really wanted was a system that could provide real-time recommendations on how to increase productivity—which crops to plant, where to grow them, how much nitrogen to use in the soil, and so on. The company developed an AI system to provide such advice, and the initial outcomes were promising; farmers were happy about the crop yields obtained with the AI's guidance. Results from that initial test were then fed back into the system to refine the algorithms used. As with the discovery step, new AI and analytic techniques can assist in co-creation by suggesting novel approaches to improving processes.

The third step for companies is to scale and then sustain the proposed solution. SEB, for example, originally deployed a version of Aida internally to assist 15,000 bank employees but thereafter rolled out the chatbot to its one million customers.

Through our work with hundreds of companies, we have identified five characteristics of business processes that companies typically want to improve: flexibility, speed, scale, decision making, and personalization. When reimagining a business process, determine which of these characteristics is central to the desired transformation, how intelligent collaboration could be harnessed to address it, and what alignments and trade-offs with other process characteristics will be necessary.

Flexibility

For Mercedes-Benz executives, inflexible processes presented a growing challenge. Increasingly, the company's most profitable customers had been demanding individualized S-class sedans, but the automaker's assembly systems couldn't deliver the customization people wanted.

Traditionally, car manufacturing has been a rigid process with automated steps executed by "dumb" robots. To improve flexibility, Mercedes replaced some of those robots with AI-enabled cobots and redesigned its processes around human-machine collaborations. At the company's plant near Stuttgart, Germany, cobot arms guided by human workers pick up and place heavy parts, becoming an extension of the worker's body.

This system puts the worker in control of the build of each car, doing less manual labor and more of a “piloting” job with the robot.

The company’s human-machine teams can adapt on the fly. In the plant, the cobots can be reprogrammed easily with a tablet, allowing them to handle different tasks depending on changes in the workflow. Such agility has enabled the manufacturer to achieve unprecedented levels of customization. Mercedes can individualize vehicle production according to the real-time choices consumers make at dealerships, changing everything from a vehicle’s dashboard components to the seat leather to the tire valve caps. As a result, no two cars rolling off the assembly line at the Stuttgart plant are the same.

Speed

For some business activities, the premium is on speed. One such operation is the detection of credit-card fraud. Companies have just seconds to determine whether they should approve a given transaction. If it’s fraudulent, they will most likely have to eat that loss. But if they deny a legitimate transaction, they lose the fee from that purchase and anger the customer.

Like most major banks, HSBC has developed an AI-based solution that improves the speed and accuracy of fraud detection. The AI monitors and scores millions of transactions daily, using data on purchase location and customer behavior, IP addresses, and other information to identify subtle patterns that signal possible fraud. HSBC first implemented the system in the

United States, significantly reducing the rate of undetected fraud and false positives, and then rolled it out in the UK and Asia. A different AI system used by Danske Bank improved its fraud-detection rate by 50% and decreased false positives by 60%. The reduction in the number of false positives frees investigators to concentrate their efforts on equivocal transactions the AI has flagged, where human judgment is needed.

The fight against financial fraud is like an arms race: Better detection leads to more-devious criminals, which leads to better detection, which continues the cycle. Thus the algorithms and scoring models for combating fraud have a very short shelf life and require continual updating. In addition, different countries and regions use different models. For these reasons, legions of data analysts, IT professionals, and experts in financial fraud are needed at the interface between humans and machines to keep the software a step ahead of the criminals.

Scale

For many business processes, poor scalability is the primary obstacle to improvement. That's particularly true of processes that depend on intensive human labor with minimal machine assistance. Consider, for instance, the employee recruitment process at Unilever. The consumer goods giant was looking for a way to diversify its 170,000-person workforce. HR determined that it needed to focus on entry-level hires and then fast-track the best into management. But the company's existing processes weren't able to evaluate potential recruits in sufficient numbers

—while giving each applicant individual attention—to ensure a diverse population of exceptional talent.

Here's how Unilever combined human and AI capabilities to scale individualized hiring: In the first round of the application process, candidates are asked to play online games that help assess traits such as risk aversion. These games have no right or wrong answers, but they help Unilever's AI figure out which individuals might be best suited for a particular position. In the next round, applicants are asked to submit a video in which they answer questions designed for the specific position they're interested in. Their responses are analyzed by an AI system that considers not just what they say but also their body language and tone. The best candidates from that round, as judged by the AI, are then invited to Unilever for in-person interviews, after which humans make the final hiring decisions.

It's too early to tell whether the new recruiting process has resulted in better employees. The company has been closely tracking the success of those hires, but more data is still needed. It is clear, however, that the new system has greatly broadened the scale of Unilever's recruiting. In part because job seekers can easily access the system by smartphone, the number of applicants doubled to 30,000 within a year, the number of universities represented surged from 840 to 2,600, and the socioeconomic diversity of new hires increased. Furthermore, the average time from application to hiring decision has dropped

from four months to just four weeks, while the time that recruiters spend reviewing applications has fallen by 75%.

Decision making

By providing employees with tailored information and guidance, AI can help them reach better decisions. This can be especially valuable for workers in the trenches, where making the right call can have a huge impact on the bottom line.

Consider the way in which equipment maintenance is being improved with the use of "digital twins"—virtual models of physical equipment. General Electric builds such software models of its turbines and other industrial products and continually updates them with operating data streaming from the equipment. By collecting readings from large numbers of machines in the field, GE has amassed a wealth of information on normal and aberrant performance. Its Predix application, which uses machine-learning algorithms, can now predict when a specific part in an individual machine might fail.

This technology has fundamentally changed the decision-intensive process of maintaining industrial equipment. Predix might, for example, identify some unexpected rotor wear and tear in a turbine, check the turbine's operational history, report that the damage has increased fourfold over the past few months, and warn that if nothing is done, the rotor will lose an estimated 70% of its useful life. The system can then suggest appropriate actions, taking into account the machine's current condition, the operating environment, and aggregated data

about similar damage and repairs to other machines. Along with its recommendations, Predix can generate information about their costs and financial benefits and provide a confidence level (say, 95%) for the assumptions used in its analysis.

Without Predix, workers would be lucky to catch the rotor damage on a routine maintenance check. It's possible that it would go undetected until the rotor failed, resulting in a costly shutdown. With Predix, maintenance workers are alerted to potential problems before they become serious, and they have the needed information at their fingertips to make good decisions—ones that can sometimes save GE millions of dollars.

Personalization

Providing customers with individually tailored brand experiences is the holy grail of marketing. With AI, such personalization can now be achieved with previously unimaginable precision and at vast scale. Think of the way the music streaming service Pandora uses AI algorithms to generate personalized playlists for each of its millions of users according to their preferences in songs, artists, and genres. Or consider Starbucks, which, with customers' permission, uses AI to recognize their mobile devices and call up their ordering history to help baristas make serving recommendations. The AI technology does what it does best, sifting through and processing copious amounts of data to recommend certain offerings or actions, and humans do what they do best,

exercising their intuition and judgment to make a recommendation or select the best fit from a set of choices.

The Carnival Corporation is applying AI to personalize the cruise experience for millions of vacationers through a wearable device called the Ocean Medallion and a network that allows smart devices to connect. Machine learning dynamically processes the data flowing from the medallion and from sensors and systems throughout the ship to help guests get the most out of their vacations. The medallion streamlines the boarding and debarking processes, tracks the guests' activities, simplifies purchasing by connecting their credit cards to the device, and acts as a room key. It also connects to a system that anticipates guests' preferences, helping crew members deliver personalized service to each guest by suggesting tailored itineraries of activities and dining experiences.

The Need for New Roles and Talent

Reimagining a business process involves more than the implementation of AI technology; it also requires a significant commitment to developing employees with what we call "fusion skills"—those that enable them to work effectively at the human-machine interface. To start, people must learn to delegate tasks to the new technology, as when physicians trust computers to help read X-rays and MRIs. Employees should also know how to combine their distinctive human skills with those of a smart machine to get a better outcome than either could

achieve alone, as in robot-assisted surgery. Workers must be able to teach intelligent agents new skills and undergo training to work well within AI-enhanced processes. For example, they must know how best to put questions to an AI agent to get the information they need. And there must be employees, like those on Apple's differential privacy team, who ensure that their companies' AI systems are used responsibly and not for illegal or unethical purposes.

We expect that in the future, company roles will be redesigned around the desired outcomes of reimagined processes, and corporations will increasingly be organized around different types of skills rather than around rigid job titles. AT&T has already begun that transition as it shifts from landline telephone services to mobile networks and starts to retrain 100,000 employees for new positions. As part of that effort, the company has completely overhauled its organizational chart: Approximately 2,000 job titles have been streamlined into a much smaller number of broad categories encompassing similar skills. Some of those skills are what one might expect (for example, proficiency in data science and data wrangling), while others are less obvious (for instance, the ability to use simple machine-learning tools to cross-sell services).

Most activities at the human-machine interface require people to *do new and different things* (such as train a chatbot) and to *do*

things differently (use that chatbot to provide better customer service). So far, however, only a small number of the companies we've surveyed have begun to reimagine their business processes to optimize collaborative intelligence. But the lesson is clear: Organizations that use machines merely to displace workers through automation will miss the full potential of AI. Such a strategy is misguided from the get-go. Tomorrow's leaders will instead be those that embrace collaborative intelligence, transforming their operations, their markets, their industries, and—no less important—their workforces.

Originally published in July–August 2018. Reprint R1804J

Stitch Fix's CEO on Selling Personal Style to the Mass Market

by Katrina Lake

AT STITCH FIX OUR BUSINESS MODEL IS SIMPLE: We send you clothing and accessories we think you'll like; you keep the items you want and send the others back. We leverage data science to deliver personalization at scale, transcending traditional brick-and-mortar and e-commerce retail experiences. Customers enjoy having an expert stylist do the shopping for them and appreciate the convenience and simplicity of the service.

Of course, making something seem simple and convenient to consumers while working profitably and at scale is complex. It's even more complex in the fashion retail industry, which is crowded, fickle, and rapidly changing. Other apparel retailers attempt to differentiate themselves through the lowest price or the fastest shipping; we differentiate ourselves through personalization. Each Fix shipment, as we call it, is a box containing five clothing and accessory items we've chosen just for you. Those choices are based on information you and

millions of others have given us—first in an extensive questionnaire you fill out when you sign up, and then in feedback you provide after each shipment.

Stitch Fix sold $730 million worth of clothing in 2016 and $977 million worth in 2017. One hundred percent of our revenue results directly from our recommendations, which are the core of our business. We have more than 2 million active clients in the United States, and we carry more than 700 brands. We're not upselling you belts that match that blouse you just added to your cart, or touting a certain brand because you've bought it before, or using browsing patterns to intuit that you might be shopping for a little black dress—all activities that have low conversion rates. Instead we make unique and personal selections by combining data and machine learning with expert human judgment.

Data science isn't woven into our culture; it *is* our culture. We started with it at the heart of the business, rather than adding it to a traditional organizational structure, and built the company's algorithms around our clients and their needs. We employ more than 80 data scientists, the majority of whom have PhDs in quantitative fields such as math, neuroscience, statistics, and astrophysics. Data science reports directly to me, and Stitch Fix wouldn't exist without data science. It's that simple.

Not a Valley Story

We're far from the prototypical Silicon Valley start-up. I don't consider myself a serial entrepreneur: Stitch Fix is the first company I've launched. But I'm fascinated by retail experiences and how untouched they were by modern technology in the 21st century. During my undergraduate years at Stanford, in the early 2000s, and in my first job, as a consultant at the Parthenon Group, I did a lot of work with retailers and restaurants. While I loved both industries and how meaningful they were to people, I was intrigued that they still provided fundamentally the same experience they had in the 1970s—or even the 1950s—despite how much the world had changed. I wondered how they might adapt, and I wanted to be part of that future.

Idea in Brief

Lake's experience as a consultant to retailers and restaurants led to a fascination with how untouched those industries were by 21st-century technology. As a lover of both clothes and data, she felt certain that data could create a better experience with apparel—as long as the human element was preserved.

From the beginning Lake planned to build a data science operation to make Stitch Fix scalable. The company's revenue is dependent on great recommendations from its algorithm, so its data scientists have a direct line to the CEO. Data science is deeply ingrained in the company culture: In addition to client recommendations of clothing, algorithms keep capital costs low, inventory moving, and deliveries efficient. Product development has adapted algorithms from genetics to find successful "traits" in clothing. Stitch Fix has even used machine learning to design apparel.

But, Lake says, shopping is inherently a personal and human activity, which is why human stylists can alter or override the product assortment a styling algorithm delivers before the client receives a shipment.

I moved on from Parthenon to become an associate at Leader Ventures, a VC firm, just as the iPhone appeared, in 2007. Still, I was thinking about retail. I studied the economics of Blockbuster during the rise of Netflix. On one side was a company that dominated physical store sales; on the other was a company that dominated sales without stores. It was the perfect case study. And I could see exactly when the scale tipped. Whenever Netflix hit about 30% market share, the local Blockbuster closed. The remaining 70% of customers then faced a decision: try Netflix or travel farther to get movies. More of them tried Netflix, putting more pressure on Blockbuster. Another store would close, and more customers would face that try-or-travel decision, in a downward spiral.

I recognized that other retailers might suffer Blockbuster's fate if they didn't rethink their strategy. For example, how would someone buy jeans 10 years down the road? I knew it wouldn't be the traditional model: go to six stores, pull pairs of jeans off the racks, try them all on. And I didn't think it would resemble today's e-commerce model either: You have 15 tabs open on your browser while you check product measurements and look for what other shoppers are saying. Then you buy multiple pairs and return the ones that don't fit.

The part of me that loves data knew it could be used to create a better experience with apparel. After all, fit and taste are just a bunch of attributes: waist, inseam, material, color, weight, durability, and pattern. It's all just data. If you collect enough, you'll get a pretty good picture of what clothes people want.

But the part of me that loves clothes recognized the human element in shopping—the feeling of finding something you weren't expecting to and delighting in the fact that it fits you and your budget. I saw an opportunity to combine those two elements—data and human experience—to create a new model for buying clothes.

A Bad Idea?

At first I didn't plan to start a company; I was going to join a start-up that wanted to pursue this idea. At Leader, I met with hundreds of entrepreneurs, hoping the right one would come through. That didn't happen. So I enrolled at Harvard Business School to pursue my risk-averse path to entrepreneurship. I used those two years to plan and launch my company. I received a term sheet to fund Stitch Fix in February 2011; I shipped the first Fix boxes from my apartment in April; and I graduated in May.

Not many people thought it was a good idea. One of my professors called it an inventory nightmare. I wanted to own all

the inventory so that I could deeply understand each item and turn it into a lot of structured data. In retail, owning all the inventory is scary, and the professor thought it would make my strategy capital-intensive and risky. But the strategy was ultimately right. Using data to better understand what people want enables us to turn over inventory faster than many conventional retailers do, because we can buy the right things and get them to the right people. Selling inventory fast enough to pay vendors with cash from clients turns out to be a very capital-efficient model.

Then there were skeptical venture capitalists. I would come to pitch meetings with a box of clothes and a personalized card from the stylist. I remember that at one meeting, a VC said within the first five minutes, "I just don't understand why anyone would ever want to receive anything like this." I appreciated his honesty. Many of them were unexcited about warehouses full of clothes. Others were baffled that we employed human stylists who were paid hourly—a very un-VC idea at a time when everything was about automation and apps. Despite our early success, Series B funding conversations got a tepid response. "I think you're great, your team is amazing, and your business is working," one VC told me. "But I get to pick one or two boards a year, and I want to pick ones I feel connected to. I can't get passionate about retail or women's dresses."

That's fair—and frustrating. As it happens, 87% of the employees, 35% of the data scientists, and 32% of the engineers at Stitch Fix are women. More than 90% of venture capitalists are men, and I felt the industry's gender dynamic was working against us. In the end, what didn't kill us made us stronger, because it forced us to focus on profitability and capital efficiency. We've since used cash from our operations to launch new businesses, including men's apparel and plus sizes for women.

Finally, there was the industry itself. By making revenue dependent on fashion recommendations, I had picked one of the more difficult tasks for machine learning. Even people who think they're undiscerning about the clothes they wear do in fact care. Fit, style, material—these matter to all of us. It's a nuanced business. That makes it especially interesting but also more difficult. Early on, focus groups asserted that they just didn't believe we could pick out clothes they'd like. They'd say, "How will it work? Nothing will fit."

The idea of paying us a $20 styling fee up front, credited to your purchase if you keep something, also gave pause. Focus group participants would ask, "Why would I pay $20 when I don't get to pick anything out?" We needed customers to trust that they'd want to keep items. And that has turned out to be true—because of the data science.

Mix and Match

STITCH FIX USES DATA that clients supply—beginning with a “style profile”—and a suite of algorithms to capture their reactions to merchandise. Human stylists (algorithmically matched with clients) review and revise every box of five items before it is mailed. Clients respond with written answers to five survey questions about each item, along with comments. That feedback, together with purchase history, allows Stitch Fix to improve its picks over time.

The following illustrates how the algorithm and the stylist together might choose one client’s very first Fix and two successive ones.

X Returned ✔ Bought

Fix 1


The client's style profile guided both the algorithm's choice of this shirt and the stylist's choice of pale pink. ✔


The stylist approved the algorithm's choice of this all-season top, even though it's out of the stated price range, because the client likes florals. ✔


These slip-on sneakers have a high match rate among clients looking for a casual shoe. The stylist thought the floral pattern would add originality. ✔


The client asked for skinny jeans. The stylist selected green from among the algorithm's denim recommendations. **X**


Because the client's style profile said she loves textures, the stylist chose this studded blouse. **X**

Fix 2


The client was looking for a versatile top. The algorithm identified this cashmere sweater because it has been extremely successful with women of her age and physical dimensions. ✔


The client did not like the fit of the green jeans, so the algorithm found a pair that fit better, and the stylist chose blue denim. **X**


The client loved the lightweight floral top in the previous box, so the stylist found this more vibrant variation, which the algorithm suggested would fit well. ✔


The client also loved the pink shirt in the previous box, so the stylist found a different take within the same color palette. ✔


The client wanted a new bag, and the algorithm found this one trending among women of her age. The stylist picked light green to pop against the red palette of the tops in the box. **X**

Fix 3


Because the client kept the cashmere sweater from the previous Fix, the stylist thought this piece, a little bolder, was worth taking a risk on. ✔


The algorithm chose this popular coat for its versatility and affordability. ✔


Stitch Fix now knows the client's preferred color and fit for jeans, so the stylist felt confident in exceeding her price range with this pair. ✔


The algorithm recommended this blouse because the client responded warmly to the color palette in the previous Fix. ✔


The stylist knows that the client is single and dating, so she chose these

 playful heels to dress up the skinny jeans. ✔

Enter the Algorithms

When I started, my "data science" was rudimentary. I used SurveyMonkey and Google Docs along with some statistical methods to track preferences and try to make good recommendations. In the beginning, I was essentially acting as a personal stylist. Sometimes I even delivered a Fix box in person. But my plan was always to build a data science operation that would make the business scalable. Our recommendations work because our algorithms are good, but our algorithms are good because data science underpins the company.

Three things make machine learning integral:

Data science reports to the CEO

At most companies, data science reports to the CTO, as part of the engineering team, or sometimes even to finance. Here it's separate, and we have a chief algorithms officer, Eric Colson, who has a seat at the strategy table. Eric came from Netflix in August 2012. Before that he was an adviser to us. He became interested in our company because it presented a challenge. At Netflix, he recalls, someone said, "What if we just started playing a movie we think someone will like when they open the app?" That seemed like a bold but risky idea—to go all in on

just one recommendation. He realized that's what Stitch Fix does. As an adviser, he found himself spending a vacation playing with some of our data. He decided to join us full-time—a huge coup for a little start-up.

Because our revenue is dependent on great recommendations from our algorithms, it's even more crucial that our data scientists have a direct line to the CEO. We also believe it sends a message to the organization as a whole about our values and our approach to strategy: Data science is extremely important, and other teams, such as marketing and engineering, will increase their capabilities by partnering closely with our data science team.

Innovation is done by data science

We've developed dozens of algorithms that no one ever asked for, because we allow our data science team to create new solutions and determine whether they have potential. No one explicitly asked the team to develop algorithms to do rebuy recommendations, for example. (Rebuys happen when a certain inventory item is selling well and we need to acquire more of it.) Our algorithms help us see these trends earlier and more accurately, so we can stock inventory more efficiently and be ready for spikes in demand. Recently the team came up with a way to track the movements of employees in our warehouses and created an algorithm that could help optimize

routes without expensive remapping of the spaces as they change.

It's sometimes hard for people to imagine how deeply ingrained data science is in our culture. We use many kinds of algorithms now, and we're building many more. Personalized recommendations of clothing, of course, are driven by machine learning. Fulfillment and inventory management use algorithms to keep capital costs low, inventory moving, and deliveries efficient. Product development has adapted some algorithms from genetics to find successful "traits" in clothes. We've even started using machine learning to design apparel.

Hybrid Designs, our in-house clothing brand, came to life one rainy afternoon when a couple of data scientists were thinking about how to fill product gaps in the marketplace. For example, many female clients in their mid-40s were asking for capped-sleeve blouses, but that style was missing from our current inventory set. Fast-forward a year, and we have 29 apparel items for women and plus sizes that were designed by computer and meet some specific, previously unfilled needs our clients have.

Another way we apply a quantitative approach to fashion is with measurement data. We track anywhere from 30 to 100 measurements on a garment, depending on what type it is, and we now know—from the experiences of more than 2 million active clients—what kind of fit would make a customer spend outside her or his comfort zone. We know the optimal ratio of

chest size to shirt width on a men's shirt. Using data analysis, we adjusted the distance from the collar to the first button on shirts for men with large chests. We know what proportion of the population fits a 27-inch inseam, and we can stock according to that proportion.

But in some ways, that's the easy part. The real challenge is having the right dress in the right color and the right size at the right time. The math around that is complex. We must account for all the measurements plus the taste of the customer, the season, the location, past trends—lots of variables.

Given a dollar to invest in the company and the choice to use it for marketing, product, or data science, we'd almost always choose data science. We're glad we started with data science at our core rather than trying to transform a traditional retailer, which I believe wouldn't have worked. For a traditional retailer to say, "Let's do what Stitch Fix does" would be like my saying, "I'd like to be taller now."

Don't forget the people

The analytical part of me loves our algorithmic approach. But shopping is inherently a personal and human activity. That's why we insist on combining data with a human stylist who can alter or override the product assortment our styling algorithm has delivered. Our stylists come from a range of design and retail backgrounds, but they all have an appreciation for the data and feel love and empathy for our clients. Humans are

much better than machines at some things—and they are likely to stay that way for a long time.

For example, when a client writes in with a very specific request, such as "I need a dress for an outdoor wedding in July," our stylists immediately know what dress options might work for that event. In addition, our clients often share intimate details of a pregnancy, a major weight loss, or a new job opportunity—all occasions whose importance a machine can't fully understand. But our stylists know exactly how special such life moments are and can go above and beyond to curate the right look, connect with the clients, and improvise when needed. That creates incredible brand loyalty.

It's simple: A good person plus a good algorithm is far superior to the best person or the best algorithm alone. We aren't pitting people and data against each other. We need them to work together. We're not training machines to behave like humans, and we're certainly not training humans to behave like machines. And we all need to acknowledge that we're fallible—the stylist, the data scientist, me. We're all wrong sometimes—even the algorithm. The important thing is that we keep learning from that.

Originally published in May–June 2018. Reprint R1803A

Strategy for Start-Ups

by Joshua Gans, Erin L. Scott, and Scott Stern

AS A START-UP, RapidSOS was an easy sell: It would bring 911 calls into the smartphone age. Emergency-response systems had evolved in a premobile era, which meant that few of them could accurately identify the location of callers who were using mobile phones, compromising response times and medical outcomes. The founders of RapidSOS—Michael Martin, an HBS graduate, and Nick Horelik, an MIT engineer—had developed a way to transmit mobile phone locations to existing 911 systems that would require only minimal adaptation on the part of other players in the emergency-services sector. After attracting early-stage financing at business plan competitions, Martin and Horelik reached a crossroads: How should they take their technology to market?

The answer wasn't straightforward—in fact, they identified four possible paths. (See the exhibit "The Entrepreneurial Strategy Compass.") They could be wildly ambitious and attempt to replace the emergency-response system altogether—creating an "Uber for ambulances." They could try a classic disruption

strategy—initially targeting poorly served populations, such as people with epilepsy, with the intention of eventually expanding to a wider swath of customers. They could avoid direct competition altogether, either by helping incumbents modernize their operations—perhaps working with 911 equipment suppliers such as Motorola—or by partnering with insurance companies, which ultimately cover the cost of ambulance service.

Many entrepreneurs, operating in the fog of uncertainty, worry that exploration will delay commercialization. They go, therefore, with the first practical strategy that comes to mind, deriding the deliberation and planning that accompany careful strategizing. As Richard Branson has famously claimed, “In the end you [have] to say, ‘Screw it, just do it’ and get on and try it.”

There are times when that approach works, of course. But usually such ad hoc experimentation should be avoided, even when it requires few resources. Entrepreneurs who commit to the first promising route they see leave their start-ups vulnerable to competitors that take a less obvious but ultimately more powerful route to commercialization and customers. Shai Agassi, for example, spent almost $1 billion building an ecosystem to support Better Place, his “swappable battery” approach to the electric car business. Elon Musk’s more deliberative, stepwise approach to developing an integrated, highly reliable Tesla turned out to be a smarter strategy.

And that’s not the only problem with an action-first philosophy. Founders are both more confident and more

persuasive to investors, employees, and partners when they can demonstrate an idea's potential across multiple strategies, validating the underlying assumptions and strength of the idea itself.

Is there a way to think through your strategic options without slowing down the process too much? After working with and studying hundreds of start-ups over the past 20 years, we have developed a framework, which we call the entrepreneurial strategy compass, that allows company founders to approach the critical choices they face in a practical and clarifying way. It delineates four generic go-to-market strategies they should consider as they move from an idea to the launch stage, each of which offers a distinct way for the venture to create and capture value.

The Entrepreneurial Strategy Compass

At the heart of our approach is the recognition that a go-to-market strategy for any innovation involves making choices about which customers to target, what technologies to apply, what organizational identity to assume, and how to position the company against which competitors. (See the sidebar "The Four Decisions.") To complicate matters, the decisions are interdependent—the choice of customers influences the company's organizational identity and its technology options.

Idea in Brief

The Problem

In their haste to get to market, entrepreneurs often run with the first plausible strategy they identify. As a result, they end up losing out to second or even third movers with superior strategies.

Why It Happens

In the innovation space it's easy to get overwhelmed by the apparent range of opportunities. Entrepreneurs fear that spending too much time weighing the alternatives will delay commercialization. The strategic commitments they make in moving forward limit their ability to pivot.

The Solution

Start-ups can improve their chances of picking the right path by investigating four generic go-to-market strategies, articulating multiple plausible versions of those strategies, and choosing the one that aligns most closely with their founders' values and motivations.

For corporations with resources, the four decisions involve analyzing data they probably already have. They can also quite often afford to engage in market research and experimentation along multiple fronts. And they can draw on prior experience. A start-up on a shoestring, in contrast, lacks a history and the knowledge it brings. However, that can actually be an advantage, because prior experience, historical data, and commitments that drive existing practices may create blind spots for established corporations, possibly even causing them to overlook innovations that pose an existential threat. Nevertheless, start-ups may ultimately face competition when incumbents wake up to new innovations, and they will definitely face pressure from other start-ups trying to beat them to market.

The Four Decisions

AT LEAST FOUR DOMAINS of decision making are crucial for every venture. Although any company will face additional choices that are particular to its context, a start-up that has not wrestled with at least these four decisions is unlikely to create and capture value on a sustainable basis. Amazon's story is illustrative.

Customers

Identifying customers and understanding their needs is usually the first step in any go-to-market strategy. But the target customer is not necessarily the first customer—and it is important that you understand the relationship between the two. You validate your product by getting the right early adopters. Amazon's decision to initially target book readers was a strategic choice. Its leadership recognized that books were a beachhead from which the company could expand into other retail categories.

Technology

Technology and customer choices are interrelated. Amazon could have built a simple online ordering system to service existing stores. Instead its goal was to let consumers buy the long tail of books that could not be stocked physically at the local mall. Thus the company had to invest beyond transaction services to build a database and a search engine capable of guiding readers through millions rather than thousands of books.

Identity, Culture, and Capabilities

Choices in this category should both create a narrative about what the company will stand for and communicate to all stakeholders what behavior to expect and what capabilities it will develop. Readers loved Amazon's offer, and Wall Street quickly saw how much money the company could make. But Amazon's founder, Jeff Bezos, wasn't building a bookstore. He wanted to create the "everything store." That would require that ordinary consumers trust they were getting a good deal,

which meant that Amazon would focus relentlessly on lowering prices, despite pressure from investors for early returns.

Competitors

Amazon defined its competition as other retailers and chose to compete aggressively by offering consumers more choice, greater reliability, and lower prices. In its early days it could easily have chosen to work with existing retailers—perhaps even defining them as customers. Competitors would have been other search and logistics service providers, and the company could have established itself as a premium service provider by adding more value for booksellers.

The entrepreneurial strategy compass

Strategic opportunities for new ventures can be categorized along two dimensions: attitude toward incumbents (collaborate or compete?) and attitude toward the innovation (build a moat or storm a hill?). This compass produces four distinct strategies that will guide a venture's decisions regarding customers, technologies, identity, and competitive space. The emergency-services provider RapidSOS used the compass to explore its strategic options.

Maintain control of the innovation and find a way to create value within the existing marketplace. Focus on being an idea factory.

For example, Dolby is the global standard setter for sound technology; it licenses proprietary technology to Sony, Bose, Apple, and others.

Create and control a new value chain, often using a platform business. Protect intellectual property.

For example, OpenTable developed a proprietary platform that allowed diners to make reservations efficiently and in so doing established influence over customer flow to restaurants.

RapidSOS could keep the technology proprietary and work with existing 911 equipment suppliers such as Motorola to modernize operations.

RapidSOS could partner with insurance companies (which ultimately pay for ambulance services); the product might take the form of a smartphone app.

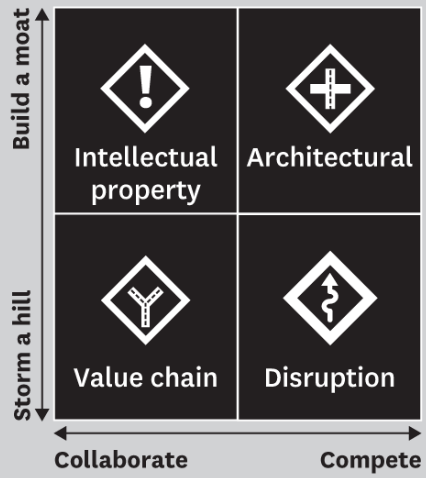

RapidSOS could replace the existing emergency response system altogether.

RapidSOS could first target poorly served populations (such as epilepsy patients) and later serve a larger swath of customers.

Focus on creating value for partners in the existing value chain. Execute quickly.

For example, Peapod became the leading U.S. internet grocer by fitting into—and improving—the grocery industry.

Compete directly with incumbents. Take them by surprise with fast execution.

For example, Rent the Runway challenged high-end retailers by offering aspiring fashion-oriented women the ability to rent rather than buy designer clothes.

Entrepreneurs may feel overwhelmed by the vast number of choices they face, even though some paths can be dismissed as impractical, and some won't coherently mesh. Our research suggests, however, that the four categories of the compass make the process manageable, getting young companies to workable go-to-market strategies quickly and laying bare the assumptions that inform choices.

To sort through potential strategies, every new venture must consider two specific competitive trade-offs:

Collaborate or compete?

Working with established players provides access to resources and supply chains that may enable the start-up to enter a larger and better-established market more quickly. Then again, the venture may encounter significant delays owing to the bureaucratic nature of large organizations and may also capture a smaller fraction of that potentially larger pie. (The incumbent is likely to hold greater bargaining power in the relationship—particularly if it can appropriate key elements of the start-up's idea.)

The alternative, too, has pluses and minuses. Competing against established players in an industry means the start-up has more freedom to build the value chain it envisions, to work with customers that the incumbents may have overlooked, and to bring innovations to market that enhance value for customers while displacing otherwise successful products. However, it means taking on competitors that have greater financial resources and an established business infrastructure.

Build a moat or storm a hill?

Some companies believe that they have more to gain from maintaining tight control over a product or a technology and that imitation will leave them vulnerable. Thus they invest in protecting intellectual property. Formal IP protection, though

expensive, can allow a technology-driven start-up to exclude others from direct competition or to wield significant bargaining power in negotiations with a supply chain partner. But prioritizing control raises the transaction costs and challenges of bringing an innovation to market and working with customers and partners.

In contrast, concentrating on quickly getting to market speeds up commercialization and development, which typically occurs in close collaboration with partners and customers. Start-ups that choose to pursue this route prioritize the ability to experiment and iterate on their ideas directly in the marketplace. Whereas a strategy built on control can delay entry, start-ups focused on getting to market expect competition and use their agility to respond when competitive threats arise. They move fast and break things.

Zeroing in on these two questions greatly simplifies the process of strategic reflection. Rather than seek to identify an á la carte combination of choices that are “right” for a given idea, a founding team can consider the potential for value creation and value capture from the various options that might be crafted within each of the four strategies.

Let’s now consider the four.

The Intellectual Property Strategy

In this quadrant of the compass, the company collaborates with incumbents and retains control of its product or technology. The

start-up focuses on idea generation and development and avoids the costs of downstream, customer-facing activities. The core idea must be of value to the customers of incumbents; therefore, development choices concerning it will dictate which incumbents are the most suitable partners for the venture.

In addition, because cooperation requires alignment with the incumbents' activities, the start-up will probably choose generalizable technology investments compatible with existing systems. Finally, the start-up's identity—as a kind of idea factory—will be reflected in its development of innovations that can be brought to market through chosen incumbents. But it will see itself as developing a small number of modular technologies that can make a decisive difference for the industry and it won't engage in unstructured experimentation with every potential new technology.

The sound company Dolby provides a quintessential example. Anyone in the market for a stereo system or watching a movie in a theater is guaranteed to come across the Dolby name. Dolby Laboratories' patented noise-reduction technologies, invented by Ray Dolby in 1965, became a global standard, retaining market leadership for 50 years. Dolby technologies have been credited with elevating the emotional intensity of iconic films such as Stanley Kubrick's *A Clockwork Orange* and George Lucas's *Star Wars.* Yet Dolby's multibillion-dollar valuation was achieved with only limited interaction with film directors, music producers, and audiophiles. The company has licensed its

proprietary technology to many product developers and manufacturers, including Sony, Bose, Apple, and Yamaha.

Entrepreneurs that pursue a strategy like Dolby's take maintaining and protecting their intellectual property very seriously. Carefully conceived patents and trademarks, managed in combination with solid R&D, can create powerful defenses that allow a start-up to preserve bargaining power over long periods of time. This strategy dictates culture and capability choices: The start-up needs to invest not only in relevant R&D skills but also in smart and committed legal minds. The IP strategy has proved powerful not only in narrow cases like Dolby's but across whole industries, such as biotechnology; with leading technology platform players, including Qualcomm; and for market intermediaries, such as Getty Images.

The Disruption Strategy

This strategy is the polar opposite of an IP strategy. It involves a decision to compete directly with incumbents, emphasizing commercialization of the idea and the rapid growth of market share rather than control of the idea's development. Disruption entrepreneurs aim to redefine established value chains and the companies that dominate those chains. But the very nature of disruption permits others to follow. Thus the heart of this strategy is the ability to get ahead and stay ahead.

Although the word "disruption" connotes chaos, the entrepreneur's initial goal is in fact to avoid poking the beast and

provoking a strong (and potentially fatal) response. The start-up strives to quickly build capabilities, resources, and customer loyalty so that when the incumbents finally wake up, the start-up is too far ahead for imitators to catch up.

For this reason, the initial choice of customers is usually a niche segment—typically one poorly served by incumbents and off their radar screen. This allows the start-up to establish credibility and explore (before anyone notices) new technologies that may have initial flaws but solid prospects for dramatic improvement. If they prove viable, these technologies are usually difficult for incumbents—whose capabilities and commitments are built around established technologies—to adopt.

The disruptive entrepreneur's identity projects hustle and verve. The start-up is staffed by the young and the hungry (and not just for ramen noodles). It doesn't fear the competitive war to come; rather, it's eager to engage. It must be lean and quick to respond. And it is intensely focused on growth.

Netflix is a poster child for this quadrant. Frustrated by movie-rental overdue fines, its founders, Marc Randolph and Reed Hastings, envisioned a solution that would leverage the then-emergent technology of DVDs. After testing their concept by sending a disc through the U.S. mail, they created a service in the late 1990s that allowed cinephiles—rather than mainstream consumers who simply wanted to watch the latest blockbuster—to receive and return DVDs that way. Netflix's strategy was to

take advantage of the “long tail” of (low-cost) content and build a recommendation engine that would reinforce customer relationships, enabling the development of a new method of movie rental that would render the brick-and-mortar Blockbuster model obsolete. (Blockbuster initially dismissed Netflix as not serving mainstream customers in a timely manner but then saw the profitability of its stores drop and ultimately disappear.)

Rent the Runway is using the disruption playbook in its drive to reshape the women’s high-end clothing market. Two Harvard MBAs, Jennifer Hyman and Jennifer Fleiss, founded the company in 2009 after identifying the challenge that fashion-oriented women faced in having to buy dresses that they might wear only once. Rent the Runway developed an online site offering aspirational women the option of renting rather than buying designer clothing and focused on solving the operational and logistical challenges of shipping dresses back and forth. Although the company has yet to displace Neiman Marcus and other more traditional players, whose focus is on wealthy haute couture customers seeking a personalized in-store experience, it has created a dedicated customer base that evangelizes the brand across social networks. Its extraordinary growth is testament to the power of execution in the face of less nimble incumbents.

The Value Chain Strategy

Disruption is exciting; by comparison, a value chain strategy seems somewhat pedestrian. The start-up invests in commercialization and day-to-day competitive strength, rather than in controlling the new product and erecting entry barriers, but its focus is on fitting into the existing value chain rather than upending it.

A pedestrian approach can nevertheless create very lucrative businesses. Consider Foxconn, the Chinese electronics manufacturer, which is one of the few global companies that can bring new products from Apple and others to market at scale and on time. The identity of such corporations arises from competence rather than aggressive competition. And although value chain entrepreneurs are driven by the customers and technology of other companies, they focus on developing scarce talent and unique capabilities to become preferred partners.

The value chain strategy is available to most start-ups. While the online grocery business Webvan, founded in 1996, was trying to disrupt the supermarket industry, Peapod became the leading U.S. internet grocer by serving as a value-added complement to traditional retailers. (Webvan went bankrupt in 2001.)

An early partnership with a Chicago-area food supplier, Jewel-Osco, allowed Peapod to clarify who its ideal customers were (professional women) and what they valued (the ability to repeat an order on a regular basis and to schedule deliveries for certain times, among other things). Whereas Webvan's disruption strategy required reconceptualizing the entire grocery-shopping

experience, Peapod's more-focused approach allowed it to develop a meaningful value proposition for customers who were willing to pay a premium for automated ordering and delivery, resulting in a profitable partnership with the supermarket chain Stop & Shop. Peapod gained the knowledge and developed the specialized capabilities with which it has led the online grocery business for nearly 20 years.

Entrepreneurs who adopt Peapod's approach create and capture value by focusing on a single "horizontal" layer of the value chain in which their expertise and capabilities are unrivaled. In probably no other entrepreneurial strategy does the founder's team play a more important role. In addition to hiring salespeople who are focused on final customers, or engineers who can improve the technical functioning of the product, it must be able to integrate innovators, business development leaders, and supply chain partners.

The start-up's capabilities must translate into enhanced differentiation or cost advantage for the established companies. And even if the innovation does enhance the competitive position of the overall value chain, the new venture can prevail only if other players in the chain are unable to replicate the value it has created.

The Architectural Strategy

Whereas the value chain strategy is the domain of quiet achievers, entrepreneurs who choose and succeed with an

architectural strategy tend to have very high public profiles. This strategy allows start-ups to both compete and achieve control, but it is out of reach for many if not most ideas and incredibly risky when it is feasible. This is the domain of Facebook and Google.

Entrepreneurs who follow an architectural strategy design an entirely new value chain and then control the key bottlenecks in it. They may not be the originators of an underlying innovation—search engines existed prior to Google, and social networks prior to Facebook—but they bring it to a mass market through careful alignment of customer, technology, and identity choices. Facebook committed early to not charging users, even though the dynamics of social media would lock them into the platform. Google adopted the motto "Don't Be Evil" so that it could achieve dominance without the pushback that had plagued other digital firms such as IBM and Microsoft. But in each case pivots were taken off the table. In other words, the risks for architectural entrepreneurs come from the fact that they may have only one shot at glory. (Remember the much-lamented Segway.)

It is perhaps not surprising that architectural entrepreneurs often end up trying to build platforms rather than products. Although platforms can be commercialized through the other strategies, if the core of a platform is closed, the entrepreneur may be able to control a new value chain.

Consider OpenTable, an online restaurant-reservation service founded in 1998 by Chuck Templeton. Motivated by the challenge of making a simple dinner reservation over the phone, Templeton hypothesized that in addition to offering a reservation platform, a successful online intermediary would have to solve the problem of restaurant-seating management. He decided to build systems that combined restaurant reservations with seating and management software, putting him in direct competition with established point-of-sale vendors such as IBM and NCR.

As Templeton recalls, OpenTable in its earliest days was “the one running wire through the rafters to get power and connectivity.” To tip the market toward his start-up, he targeted the most influential restaurants first. “We were able to get the top 20 restaurants [in San Francisco],” he says, “and the next 50 would all want to be where those top 20 were. There began to be a critical mass on the website.” Templeton reorganized the value chain of the dining industry so that the internal operations of restaurants were integrated into customers’ first engagement with them: the reservation phase. OpenTable achieved control over valuable proprietary data on customer preferences and demand and established a hard-to-dislodge platform that is “table stakes” for a new restaurateur. This dominance underlay its $2.6 billion acquisition by Priceline in 2014.

Let’s look now at how entrepreneurs can use the strategy compass to decide among the four basic approaches.

Making the Choice

The first step is to fill as many of the quadrants of the compass as possible with strategic options. This is no simple task. It involves gathering additional information and experimenting to some degree (but commitments should be modest until a choice is made).

Particularly effective approaches for start-ups can be found in Eric Ries's *The Lean Startup*, Alexander Osterwalder and Yves Pigneur's *Business Model Generation,* and Bill Aulet's *Disciplined Entrepreneurship*. Whatever framework is chosen, however, it should involve an explicit process of hypothesis building and testing—an observation that was nicely made in "Bringing Science to the Art of Strategy," by A.G. Lafley, Roger L. Martin, Jan W. Rivkin, and Nicolaj Siggelkow (HBR, September 2012).

This process at a minimum yields crucial insight into stumbling blocks associated with particular paths within the compass. Some alternatives can be dismissed owing to lack of feasibility or lack of alignment with the capabilities of the founding team. In other cases, the requirements—in terms of capital, commitment, and momentum—will be clear, allowing the start-up to focus on them to make the chosen strategy work.

Once the alternatives have been identified, how should the entrepreneur actually make a choice? Let's go back to RapidSOS. As the founders debated the next steps for their idea—mobile-centric emergency-response systems—they used the compass to identify four strategies. As noted earlier, they could use an

architectural strategy to replace the existing 911 system with an "Uber for ambulances." They could use an IP strategy to collaborate with existing players in the emergency-response sector. They could use a value chain strategy to work with insurance companies and other consumer-facing partners, becoming a feature for a corporate smartphone app. Or they could use a disruption strategy to focus on a narrow customer segment for whom emergency response is a priority—such as epileptics—and partner with patient advocacy groups to meet its needs.

For each compass quadrant the company identified which customers to target, which technologies to focus on, what identity to assume, and whom to compete with and how. All four paths looked plausible, which was a striking validation of the founders' idea. If only one viable vision of the future exists, the entrepreneur probably doesn't have much of a business to begin with.

Having several good options need not be paralyzing. Quite simply, entrepreneurs should choose the strategy that aligns best with the purpose they originally brought to the venture. The RapidSOS mission to improve services for specific patient groups led the team to focus with a high level of conviction on a disruption strategy. This commitment—which Martin and Horelik could communicate with passion and purpose—allowed them to win over patient groups and stakeholders throughout

the emergency-response sector, enabling RapidSOS to roll out its technology to the broader market over two years.

The founding team does not just make the choice; it has to live the choice. Alignment between strategy and purpose is crucial for motivating founders and persuading early stakeholders to travel the chosen path. To be clear, making a choice requires commitment but does not foreclose all other paths forward. RapidSOS's decision to engage with both patient advocates and the emergency-response community meant that the start-up was unlikely to bypass traditional 911 systems—at least in the medium term. But the focus on patient advocacy groups encouraged end-user engagement, which over time generated meaningful collaboration opportunities and attracted investment from more-established players, including Motorola.

Still, every strategy affects possible future pivots, removing some and opening up others. A venture must be mindful of this so that it doesn't raise future costs but does enable opportunities to move from the start-up to the scale-up phase.

The entrepreneurial strategy compass does not eliminate or minimize the uncertainty inherent in launching a start-up. What it does is provide a coherent framework for escaping the perceived realities of the existing environment and defining possible new environments to choose from. The word "choose" is critical here: When a start-up is competing with new products

in the absence of a significant innovation, its success is largely determined by how its strategic choices are informed by the environment. Among established businesses, the winner is usually the company that understands the environment better. But entrepreneurs offering something significantly new have an opportunity to reshape the environment—perhaps, as with Dolby, to create a part of it that they will own or, as with Amazon, to create an altogether different reality. Which they choose is largely up to them. Our framework is designed to help them make that choice successfully and channel imagination and commitment toward the realization of their ideas.

Originally printed in May–June 2018. Reprint R1803B

Erin L. Scott and Scott Stern gratefully acknowledge financial support from the Jean Hammond (1986) and Michael Krasner (1974) Entrepreneurship Fund, the Edward B. Roberts (1957) Entrepreneurship Fund, and the MIT Sloan School of Management.

Agile at Scale

by Darrell K. Rigby, Jeff Sutherland, and Andy Noble

BY NOW MOST BUSINESS LEADERS are familiar with agile innovation teams. These small, entrepreneurial groups are designed to stay close to customers and adapt quickly to changing conditions. When implemented correctly, they almost always result in higher team productivity and morale, faster time to market, better quality, and lower risk than traditional approaches can achieve.

Naturally, leaders who have experienced or heard about agile teams are asking some compelling questions. What if a company were to launch dozens, hundreds, or even thousands of agile teams throughout the organization? Could whole segments of the business learn to operate in this manner? Would scaling up agile improve corporate performance as much as agile methods improve individual team performance?

In today's tumultuous markets, where established companies are furiously battling assaults from startups and other insurgent competitors, the prospect of a fast-moving, adaptive

organization is highly appealing. But as enticing as such a vision is, turning it into a reality can be challenging. Companies often struggle to know which functions should be reorganized into multidisciplinary agile teams and which should not. And it's not unusual to launch hundreds of new agile teams only to see them bottlenecked by slow-moving bureaucracies.

We have studied the scaling up of agile at hundreds of companies, including small firms that run the entire enterprise with agile methods; larger companies that, like Spotify and Netflix, were born agile and have become more so as they've grown; and companies that, like Amazon and USAA (the financial services company for the military community), are making the transition from traditional hierarchies to more-agile enterprises. Along with the many success stories are some disappointments. For example, one prominent industrial company's attempts over the past five years to innovate like a lean start-up have not yet generated the financial results sought by activist investors and the board of directors, and several senior executives recently resigned.

Our studies show that companies can scale up agile effectively and that doing so creates substantial benefits. But leaders must be realistic. Not every function needs to be organized into agile teams; indeed, agile methods aren't well suited to some activities. Once you begin launching dozens or hundreds of agile teams, however, you can't just leave the other parts of the business alone. If your newly agile units are

constantly frustrated by bureaucratic procedures or a lack of collaboration between operations and innovation teams, sparks will fly from the organizational friction, leading to meltdowns and poor results. Changes are necessary to ensure that the functions that don't operate as agile teams support the ones that do.

Leading Agile by Being Agile

For anyone who isn't familiar with agile, here's a short review. Agile teams are best suited to innovation—that is, the profitable application of creativity to improve products and services, processes, or business models. They are small and multidisciplinary. Confronted with a large, complex problem, they break it into modules, develop solutions to each component through rapid prototyping and tight feedback loops, and integrate the solutions into a coherent whole. They place more value on adapting to change than on sticking to a plan, and they hold themselves accountable for outcomes (such as growth, profitability, and customer loyalty), not outputs (such as lines of code or number of new products).

Idea in Brief

The Ambition

To go from a handful of agile innovation teams in a function like software development to scores, even hundreds, throughout your company—to make agile the dominant way you operate.

The Challenges

Figuring out where to start and how fast and far to go, deciding which functions can and should be converted to agile teams and which should not, and preventing slow-moving bureaucracies from impeding those that do convert.

The Solution

Leaders should use agile methods themselves and create a *taxonomy of opportunities* to set priorities and break the journey into small steps. Workstreams should be modularized and then seamlessly integrated. Functions not reorganized into agile teams should learn to operate with agile values. The annual budgeting process should be complemented with a VC-like approach to funding.

Conditions are ripe for agile teams in any situation where problems are complex, solutions are at first unclear, project requirements are likely to change, close collaboration with end users is feasible, and creative teams will outperform command-and-control groups. Routine operations such as plant maintenance, purchasing, and accounting are less fertile ground. Agile methods caught on first in IT departments and are now widely used in software development. Over time they have spread into functions such as product development, marketing, and even HR. (See "Embracing Agile," HBR, May 2016, and "HR Goes Agile," HBR, March–April 2018.)

Agile teams work differently from chain-of-command bureaucracies. They are largely self-governing: Senior leaders tell team members where to innovate but not how. And the teams work closely with customers, both external and internal.

Ideally, this puts responsibility for innovation in the hands of those who are closest to customers. It reduces layers of control and approval, thereby speeding up work and increasing the teams' motivation. It also frees up senior leaders to do what only they can do: create and communicate long-term visions, set and sequence strategic priorities, and build the organizational capabilities to achieve those goals.

When leaders haven't themselves understood and adopted agile approaches, they may try to scale up agile the way they have attacked other change initiatives: through top-down plans and directives. The track record is better when they behave like an agile team. That means viewing various parts of the organization as their customers—people and groups whose needs differ, are probably misunderstood, and will evolve as agile takes hold. The executive team sets priorities and sequences opportunities to improve those customers' experiences and increase their success. Leaders plunge in to solve problems and remove constraints rather than delegate that work to subordinates. The agile leadership team, like any other agile team, has an "initiative owner" who is responsible for overall results and a facilitator who coaches team members and helps keep everyone actively engaged.

Bosch, a leading global supplier of technology and services with more than 400,000 associates and operations in 60-plus countries, took this approach. As leaders began to see that traditional top-down management was no longer effective in a

fast-moving, globalized world, the company became an early adopter of agile methods. But different business areas required different approaches, and Bosch's first attempt to implement what it called a "dual organization"—one in which hot new businesses were run with agile teams while traditional functions were left out of the action—compromised the goal of a holistic transformation. In 2015 members of the board of management, led by CEO Volkmar Denner, decided to build a more unified approach to agile teams. The board acted as a steering committee and named Felix Hieronymi, a software engineer turned agile expert, to guide the effort.

At first Hieronymi expected to manage the assignment the same way Bosch managed most projects: with a goal, a target completion date, and regular status reports to the board. But that approach felt inconsistent with agile principles, and the company's divisions were just too skeptical of yet another centrally organized program. So the team shifted gears. "The steering committee turned into a working committee," Hieronymi told us. "The discussions got far more interactive." The team compiled and rank-ordered a backlog of corporate priorities that was regularly updated, and it focused on steadily removing companywide barriers to greater agility. Members fanned out to engage division leaders in dialogue. "Strategy evolved from an annual project to a continuous process," Hieronymi says. "The members of the management board divided themselves into small agile teams and tested various

approaches—some with a ‘product owner’ and an ‘agile master’—to tackle tough problems or work on fundamental topics. One group, for instance, drafted the 10 new leadership principles released in 2016. They personally experienced the satisfaction of increasing speed and effectiveness. You can’t gain this experience by reading a book.” Today Bosch operates with a mix of agile teams and traditionally structured units. But it reports that nearly all areas have adopted agile values, are collaborating more effectively, and are adapting more quickly to increasingly dynamic marketplaces.

Getting Agile Rolling

At Bosch and other advanced agile enterprises, the visions are ambitious. In keeping with agile principles, however, the leadership team doesn’t plan every detail in advance. Leaders recognize that they do not yet know how many agile teams they will require, how quickly they should add them, and how they can address bureaucratic constraints without throwing the organization into chaos. So they typically launch an initial wave of agile teams, gather data on the value those teams create and the constraints they face, and then decide whether, when, and how to take the next step. This lets them weigh the value of increasing agility (in terms of financial results, customer outcomes, and employee performance) against its costs (in terms of both financial investments and organizational

challenges). If the benefits outweigh the costs, leaders continue to scale up agile—deploying another wave of teams, unblocking constraints in less agile parts of the organization, and repeating the cycle. If not, they can pause, monitor the market environment, and explore ways to increase the value of the agile teams already in place (for instance, by improving the prioritization of work or upgrading prototyping capabilities) and decrease the costs of change (by publicizing agile successes or hiring experienced agile enthusiasts).

To get started on this test-and-learn cycle, leadership teams typically employ two essential tools: a taxonomy of potential teams and a sequencing plan reflecting the company's key priorities. Let's first look at how each can be employed and then explore what more is needed to tackle large-scale, long-term agile initiatives.

Create a taxonomy of teams

Just as agile teams compile a backlog of work to be accomplished in the future, companies that successfully scale up agile usually begin by creating a full taxonomy of opportunities. Following agile's modular approach, they may break the taxonomy into three components—customer experience teams, business process teams, and technology systems teams—and then integrate them. The first component identifies all the experiences that could significantly affect external and internal customer decisions, behaviors, and

satisfaction. These can usually be divided into a dozen or so major experiences (for example, one of a retail customer's major experiences is to buy and pay for a product), which in turn can be divided into dozens of more-specific experiences (the customer may need to choose a payment method, use a coupon, redeem loyalty points, complete the checkout process, and get a receipt). The second component examines the relationships among these experiences and key business processes (improved checkout to reduce time in lines, for instance), aiming to reduce overlapping responsibilities and increase collaboration between process teams and customer experience teams. The third focuses on developing technology systems (such as better mobile-checkout apps) to improve the processes that will support customer experience teams.

The taxonomy of a $10 billion business might identify anywhere from 350 to 1,000 or more potential teams. Those numbers sound daunting, and senior executives are often loath even to consider so much change ("How about if we try two or three of these things and see how it goes?"). But the value of a taxonomy is that it encourages exploration of a transformational vision while breaking the journey into small steps that can be paused, turned, or halted at any time. It also helps leaders spot constraints. Once you've identified the teams you could launch and the sorts of people you would need to staff them, for instance, you need to ask: Do we have those people? If so, where are they? A taxonomy reveals your

talent gaps and the kinds of people you must hire or retrain to fill them. Leaders can also see how each potential team fits into the goal of delivering better customer experiences.

USAA has more than 500 agile teams up and running and plans to add 100 more in 2018. The taxonomy is fully visible to everyone across the enterprise. "If you don't have a really good taxonomy, you get redundancy and duplication," COO Carl Liebert told us. "I want to walk into an auditorium and ask, 'Who owns the member's change-of-address experience?' And I want a clear and confident response from a team that owns that experience, whether a member is calling us, logging into our website on a laptop, or using our mobile app. No finger-pointing. No answers that begin with 'It's complicated.'?"

USAA's taxonomy ties the activities of agile teams to the people responsible for business units and product lines. The goal is to ensure that managers responsible for specific parts of the P&L understand how cross-functional teams will influence their results. The company has senior leaders who act as general managers in each line of business and are fully accountable for business results. But those leaders rely on customer-focused, cross-organizational teams to get much of the work done. The company also depends on technology and digital resources assigned to the experience owners; the goal here is to ensure that business leaders have the end-to-end resources to deliver the outcomes they have committed to. The intent of the taxonomy is to clarify how to engage the right

people in the right work without creating confusion. This kind of link is especially important when hierarchical organizational structures do not align with customer behaviors. For example, many companies have separate structures and P&Ls for online and offline operations—but customers want seamlessly integrated omnichannel experiences. A clear taxonomy that launches the right cross-organizational teams makes such alignment possible.

Sequence the transition

Taxonomy in hand, the leadership team sets priorities and sequences initiatives. Leaders must consider multiple criteria, including strategic importance, budget limitations, availability of people, return on investment, cost of delays, risk levels, and interdependencies among teams. The most important—and the most frequently overlooked—are the pain points felt by customers and employees on the one hand and the organization's capabilities and constraints on the other. These determine the right balance between how fast the rollout should proceed and how many teams the organization can handle simultaneously.

A few companies, facing urgent strategic threats and in need of radical change, have pursued big-bang, everything-at-once deployments in some units. For example, in 2015 ING Netherlands anticipated rising customer demand for digital solutions and increasing incursions by new digital competitors

(“fintechs”). The management team decided to move aggressively. It dissolved the organizational structures of its most innovative functions, including IT development, product management, channel management, and marketing—essentially abolishing everyone’s job. Then it created small agile “squads” and required nearly 3,500 employees to reapply for 2,500 redesigned positions on those squads. About 40% of the people filling the positions had to learn new jobs, and all had to profoundly change their mindset. (See “One Bank’s Agile Team Experiment,” HBR, March–April 2018.)

But big-bang transitions are hard. They require total leadership commitment, a receptive culture, enough talented and experienced agile practitioners to staff hundreds of teams without depleting other capabilities, and highly prescriptive instruction manuals to align everyone’s approach. They also require a high tolerance of risk, along with contingency plans to deal with unexpected breakdowns. ING continues to iron out wrinkles as it expands agile throughout the organization.

Companies short on those assets are better off rolling out agile in sequenced steps, with each unit matching the implementation of opportunities to its capabilities. At the beginning of its agile initiative, the advanced technology group at 3M Health Information Systems launched eight to 10 teams every month or two; now, two years in, more than 90 teams are up and running. 3M’s Corporate Research Systems Lab got started later but launched 20 teams in three months.

Whatever the pace or endpoint, results should begin showing up quickly. Financial results may take a while—Jeff Bezos believes that most initiatives take five to seven years to pay dividends for Amazon—but positive changes in customer behavior and team problem solving provide early signs that initiatives are on the right track. “Agile adoption has already enabled accelerated product deliveries and the release of a beta application six months earlier than originally planned,” says Tammy Sparrow, a senior program manager at 3M Health Information Systems.

Division leaders can determine the sequencing just as any agile team would. Start with the initiatives that offer potentially the greatest value and the most learning. SAP, the enterprise software company, was an early scaler of agile, launching the process a decade ago. Its leaders expanded agile first in its software development units—a highly customer-centric segment where they could test and refine the approach. They established a small consulting group to train, coach, and embed the new way of working, and they created a results tracker so that everyone could see the teams’ gains. “Showing concrete examples of impressive productivity gains from agile created more and more pull from the organization,” says Sebastian Wagner, who was then a consulting manager in that group. Over the next two years the company rolled out agile to more than 80% of its development organizations, creating more than 2,000 teams. People in sales and marketing saw the

need to adapt in order to keep up, so those areas went next. Once the front end of the business was moving at speed, it was time for the back end to make the leap, so SAP shifted its group working on internal IT systems to agile.

Too many companies make the mistake of going for easy wins. They put teams into offsite incubators. They intervene to create easy workarounds to systemic obstacles. Such coddling increases the odds of a team's success, but it doesn't produce the learning environment or organizational changes necessary to scale dozens or hundreds of teams. A company's early agile teams carry the burden of destiny. Testing them, just like testing any prototype, should reflect diverse, realistic conditions. Like SAP, the most successful companies focus on vital customer experiences that cause the greatest frustrations among functional silos.

Still, no agile team should launch unless and until it is ready to begin. *Ready* doesn't mean planned in detail and guaranteed to succeed. It means that the team is:

- focused on a major business opportunity with a lot at stake
- responsible for specific outcomes
- trusted to work autonomously—guided by clear decision rights, properly resourced, and staffed with a small group of multidisciplinary experts who are passionate about the opportunity

- committed to applying agile values, principles, and practices
- empowered to collaborate closely with customers
- able to create rapid prototypes and fast feedback loops
- supported by senior executives who will address impediments and drive adoption of the team's work

Following this checklist will help you plot your sequence for the greatest impact on both customers and the organization.

Master large-scale agile initiatives

Many executives have trouble imagining that small agile teams can attack large-scale, long-term projects. But in principle there is no limit to the number of agile teams you can create or how large the initiative can be. You can establish "teams of teams" that work on related initiatives—an approach that is highly scalable. Saab's aeronautics business, for instance, has more than 100 agile teams operating across software, hardware, and fuselage for its Gripen fighter jet—a $43 million item that is certainly one of the most complex products in the world. It coordinates through daily team-of-teams stand-ups. At 7:30 AM each frontline agile team holds a 15-minute meeting to flag impediments, some of which cannot be resolved within that team. At 7:45 the impediments requiring coordination are escalated to a team of teams, where leaders work to either settle or further escalate issues. This approach continues, and

by 8:45 the executive action team has a list of the critical issues it must resolve to keep progress on track. Aeronautics also coordinates its teams through a common rhythm of three-week sprints, a project master plan that is treated as a living document, and the colocation of traditionally disparate parts of the organization—for instance, putting test pilots and simulators with development teams. The results are dramatic: IHS Jane's has deemed the Gripen the world's most cost-effective military aircraft.

Building Agility Across the Business

Expanding the number of agile teams is an important step toward increasing the agility of a business. But equally important is how those teams interact with the rest of the organization. Even the most advanced agile enterprises—Amazon, Spotify, Google, Netflix, Bosch, Saab, SAP, Salesforce, Riot Games, Tesla, and SpaceX, to name a few—operate with a mix of agile teams and traditional structures. To ensure that bureaucratic functions don't hamper the work of agile teams or fail to adopt and commercialize the innovations developed by those teams, such companies constantly push for greater change in at least four areas.

Values and principles

A traditional hierarchical company can usually accommodate a small number of agile teams sprinkled around the organization.

Conflicts between the teams and conventional procedures can be resolved through personal interventions and workarounds. When a company launches several hundred agile teams, however, that kind of ad hoc accommodation is no longer possible. Agile teams will be pressing ahead on every front. Traditionally structured parts of the organization will fiercely defend the status quo. As with any change, skeptics can and will produce all kinds of antibodies that attack agile, ranging from refusals to operate on an agile timetable (“Sorry, we can’t get to that software module you need for six months”) to the withholding of funds from big opportunities that require unfamiliar solutions.

So a leadership team hoping to scale up agile needs to instill agile values and principles throughout the enterprise, including the parts that do not organize into agile teams. This is why Bosch’s leaders developed new leadership principles and fanned out throughout the company: They wanted to ensure that everyone understood that things would be different and that agile would be at the center of the company’s culture.

Operating architectures

Implementing agile at scale requires modularizing and then seamlessly integrating workstreams. For example, Amazon can deploy software thousands of times a day because its IT architecture was designed to help developers make fast, frequent releases without jeopardizing the firm’s complex

systems. But many large companies, no matter how fast they can code programs, can deploy software only a few times a day or a week; that's how their architecture works.

Building on the modular approach to product development pioneered by Toyota, Tesla meticulously designs interfaces among the components of its cars to allow each module to innovate independently. Thus the bumper team can change anything as long as it maintains stable interfaces with the parts it affects. Tesla is also abandoning traditional annual release cycles in favor of real-time responses to customer feedback. CEO Elon Musk says that the company makes about 20 engineering changes a week to improve the production and performance of the Model S. Examples include new battery packs, updated safety and autopilot hardware, and software that automatically adjusts the steering wheel and seat for easier entry and exit.

In the most advanced agile enterprises, innovative product and process architectures are attacking some of the thorniest organizational constraints to further scaling. Riot Games, the developer of the wildly successful multiplayer online battle arena League of Legends, is redesigning the interfaces between agile teams and support-and-control functions that operate conventionally, such as facilities, finance, and HR. Brandon Hsiung, the product lead for this ongoing initiative, says it involves at least two key steps. One is shifting the functions' definition of their customers. "Their customers are not their

functional bosses, or the CEO, or even the board of directors," he explains. "Their customers are the development teams they serve, who ultimately serve our players." The company instituted Net Promoter surveys to collect feedback on whether those customers would recommend the functions to others and made it plain that dissatisfied customers could sometimes hire outside providers. "It's the last thing we want to happen, but we want to make sure our functions develop world-class capabilities that could compete in a free market," Hsiung says.

Riot Games also revamped how its corporate functions interact with its agile teams. Some members of corporate functions may be embedded in agile teams, or a portion of a function's capacity may be dedicated to requests from agile teams. Alternatively, functions might have little formal engagement with the teams after collaborating with them to establish certain boundaries. Says Hsiung: "Silos such as real estate and learning and development might publish philosophies, guidelines, and rules and then say, 'Here are our guidelines. As long as you operate within them, you can go crazy; do whatever you believe is best for our players.'"

In companies that have scaled up agile, the organization charts of support functions and routine operations generally look much as they did before, though often with fewer management layers and broader spans of control as supervisors learn to trust and empower people. The bigger changes are in the ways functional departments work. Functional priorities

are necessarily more fully aligned with corporate strategies. If one of the company's key priorities is improving customers' mobile experience, that can't be number 15 on finance's funding list or HR's hiring list. And departments such as legal may need buffer capacity to deal with urgent requests from high-priority agile teams.

Over time even routine operations with hierarchical structures are likely to develop more-agile mindsets. Of course, finance departments will always manage budgets, but they don't need to keep questioning the decisions of the owners of agile initiatives. "Our CFO constantly shifts accountability to empowered agile teams," says Ahmed Sidky, the head of development management at Riot Games. "He'll say, 'I am not here to run the finances of the company. You are, as team leaders. I'm here in an advisory capacity.' In the day-to-day organization, finance partners are embedded in every team. They don't control what the teams do or don't do. They are more like finance coaches who ask hard questions and provide deep expertise. But ultimately it's the team leader who makes decisions, according to what is best for Riot players."

Some companies, and some individuals, may find these trade-offs hard to accept and challenging to implement. Reducing control is always scary—until you do so and find that people are happier and success rates triple. In a recent Bain survey of nearly 1,300 global executives, more respondents agreed with this statement about management than with any

other: “Today’s business leaders must trust and empower people, not command and control them.” (Only 5% disagreed.)

Talent acquisition and motivation

Companies that are scaling up agile need systems for acquiring star players and motivating them to make teams better. (Treat your stars unfairly, and they will bolt to a sexy start-up.) They also need to unleash the wasted potential of more-typical team members and build commitment, trust, and joint accountability for outcomes. There’s no practical way to do this without changing HR procedures. A company can no longer hire purely for expertise, for instance; it now needs expertise combined with enthusiasm for work on a collaborative team. It can’t evaluate people according to whether they hit individual objectives; it now needs to look at their performance on agile teams and at team members’ evaluations of one another. Performance assessments typically shift from an annual basis to a system that provides relevant feedback and coaching every few weeks or months. Training and coaching programs encourage the development of cross-functional skills customized to the needs of individual employees. Job titles matter less and change less frequently with self-governing teams and fewer hierarchical levels. Career paths show how product owners—the individuals who set the vision and own the results of an agile team—can continue their personal

development, expand their influence, and increase their compensation.

Companies may also need to revamp their compensation systems to reward group rather than individual accomplishments. They need recognition programs that celebrate contributions immediately. Public recognition is better than confidential cash bonuses at bolstering agile values—it inspires recipients to improve even further, and it motivates others to emulate the recipients' behaviors. Leaders can also reward "A" players by engaging them in the most vital opportunities, providing them with the most advanced tools and the greatest possible freedom, and connecting them with the most talented mentors in their field.

Annual planning and budgeting cycles

In bureaucratic companies, annual strategy sessions and budget negotiations are powerful tools for aligning the organization and securing commitments to stretch goals. Agile practitioners begin with different assumptions. They see that customer needs change frequently and that breakthrough insights can occur at any time. In their view, annual cycles constrain innovation and adaptation: Unproductive projects burn resources until their budgets run out, while critical innovations wait in line for the next budget cycle to compete for funding.

In companies with many agile teams, funding procedures are different. Funders recognize that for two-thirds of successful innovations, the original concept will change significantly during the development process. They expect that teams will drop some features and launch others without waiting for the next annual cycle. As a result, funding procedures evolve to resemble those of a venture capitalist. VCs typically view funding decisions as opportunities to purchase options for further discovery. The objective is not to instantly create a large-scale business but, rather, to find a critical component of the ultimate solution. This leads to a lot of apparent failures but accelerates and reduces the cost of learning. Such an approach works well in an agile enterprise, vastly improving the speed and efficiency of innovation.

Companies that successfully scale up agile see major changes in their business. Scaling up shifts the mix of work so that the business is doing more innovation relative to routine operations. The business is better able to read changing conditions and priorities, develop adaptive solutions, and avoid the constant crises that so frequently hit traditional hierarchies. Disruptive innovations will come to feel less disruptive and more like adaptive business as usual. The scaling up also brings agile values and principles to business operations and support functions, even if many routine

activities remain. It leads to greater efficiency and productivity in some of the business's big cost centers. It improves operating architectures and organizational models to enhance coordination between agile teams and routine operations. Changes come on line faster and are more responsive to customer needs. Finally, the business delivers measurable improvements in outcomes—not only better financial results but also greater customer loyalty and employee engagement.

Agile's test-and-learn approach is often described as incremental and iterative, but no one should mistake incremental development processes for incremental thinking. SpaceX, for example, aims to use agile innovation to begin transporting people to Mars by 2024, with the goal of establishing a self-sustaining colony on the planet. How will that happen? Well, people at the company don't really know . . . yet. But they have a vision that it's possible, and they have some steps in mind. They intend to dramatically improve reliability and reduce expenses, partly by reusing rockets much like airplanes. They intend to improve propulsion systems to launch rockets that can carry at least 100 people. They plan to figure out how to refuel in space. Some of the steps include pushing current technologies as far as possible and then waiting for new partners and new technologies to emerge.

That's agile in practice: big ambitions and step-by-step progress. It shows the way to proceed even when, as is so often the case, the future is murky.

Originally published in May–June 2018. Reprint R1803F

Operational Transparency

by Ryan W. Buell

BARCLAYS BANK INSTALLED the world's first successful automated teller machine to much fanfare in June 1967. Having a machine distribute cash was less expensive and more efficient than having a human teller do it. What's more, customers could access the ATM at any hour—even when the bank was closed. It seemed like a win-win, and ATMs quickly spread around the world. Today people are three times more likely to withdraw money from an ATM than from a human teller.

However, there's a wrinkle to the ATM success story. When customers use ATMs more and tellers less, their overall level of satisfaction with their bank goes down. It turns out that when consumers can't see the work that's being done to serve them, their perception is that less effort went into delivering the service, so they don't appreciate or value it as much. ATMs carry out complex work: They reliably identify customers, find their account information, and then accurately complete the transaction—all while protecting the confidentiality of their

private information. But separated from this work by a hard, metallic surface and a vague “processing transaction” message, customers take the “wizardry” for granted in a way that they don’t when they’re face-to-face with tellers who are working in their behalf.

Automation has enabled enormous efficiencies in recent years, but it has also detached customers from operations. Thanks to fleets of order-picking robots and miles of automated conveyors, it takes less than one minute of human labor to pick, pack, and ship the typical Amazon package—a miraculous ballet among people and machines that customers never glimpse. Google has more than a million servers working to deliver answers to more than a trillion queries a year—information distributed in fractions of a second without a hint of the massive operation behind it.

And even where technology hasn’t erected barriers between customers and the work being performed for them, leaders have put them up. At hospitals, as many as 70% of clinical diagnoses come from the pathology lab—but the people who run those tests are often hidden away in the basement or off-site. Hundreds of people have a hand in the successful takeoff and landing of a commercial flight—but for the most part, passengers see only the cabin crew. Consider all the people who work in offices, kitchens, warehouses, and factories whose efforts create immeasurable value but who never enter customers’ minds.

Therein lies a crucial managerial dilemma that I've been studying over the past decade. It has long been believed that the more contact an operation has with its customers, the less efficiently it runs. Customers are, as a researcher in the 1960s boldly called them, "environmental disturbances." As the argument goes, separating customers from internal processes through physical distance, time, or the introduction of technology enables companies to perform more efficiently and, in turn, create more value for consumers. But my research shows that the pendulum can swing too far. When customers are cordoned off from a company's operation, they are less likely to fully understand and appreciate the value being created. As a result, they are less satisfied, less willing to pay, less trusting, and less loyal to the company over time. Employees also suffer when they are cut off from the business's front lines, as they lose the motivation and enjoyment that comes from making a difference in people's lives and are denied the opportunities to learn and improve that arise from interaction with customers.

Idea in Brief

The Dilemma

Conventional wisdom holds that the more contact an operation has with its customers, the less efficiently it will run. But when customers are partitioned away from the operation, they are less likely to fully understand and appreciate the work going on behind the scenes,

thereby placing a lower value on the product or service being offered.

The Solution

Managers should experiment with operational transparency—the deliberate design of windows into and out of the organization's operations to help customers understand and appreciate the value being added.

The Benefits

Witnessing the hidden work performed on their behalf makes customers more satisfied, more willing to pay, and more loyal. It can also make employees more satisfied by demonstrating to them that they are serving their customers well. However, managers should be aware of certain conditions in which transparency can backfire.

One solution that my colleagues and I have investigated is the introduction of operational transparency—the deliberate design of windows into and out of the organization's operations to help customers and employees alike understand and appreciate the value being created. To determine when and how to design such windows, managers must understand when and how customers and employees want to open up operations to scrutiny—and when both parties would prefer that work be undertaken behind the scenes.

Behind the Curtain

I first started documenting the beneficial effects of operational transparency in 2008, when I set up a mock website called

Travel Finder, with my Harvard Business School colleague Michael Norton, as part of a study. We had noticed that travel agents, like bank tellers, were being made increasingly obsolete by technology—in this case, by online travel agencies. We also noticed that most online ticket sites hid the work they performed for customers behind progress bars and activity spinners, or behind marketing messages such as “Did you know you can book your hotel with us, too?” Online travel agency Kayak was an exception. The company showed customers how many different airlines it was searching while they waited, and it slotted itineraries into the results screen as they were found instead of all at once. We wondered whether this operational transparency would change the way customers viewed the service.

For our travel study, we recruited people to search for flights from Boston to Los Angeles on our website. After they entered their search information, we randomly varied how long people waited as the website searched for possible tickets. While waiting, some people saw a progress bar, and some were shown, in addition to the progress bar, the hidden work that the website was doing: “Now getting results from American Airlines . . . from JetBlue . . . 133 results found so far . . . Now 427 . . .” We then surveyed people about how valuable they perceived the website to be. No matter how long people had waited, they always considered the website to be more valuable when it showed the work it was doing for them. They

also reported a higher willingness to pay, a perception of higher quality, and a greater desire to use the site again. What's more, they were also considerably less sensitive to their wait time when they experienced operational transparency. People who received instantaneous service perceived the service to be as valuable as people who waited 25 seconds with a progress bar, and as valuable as people who waited 55 seconds with operational transparency. That's remarkable in an era in which we have come to expect online services to be delivered in fractions of a second.

In other experiments, people who experienced operational transparency expressed more interest in using the website again in the future, even when they compared it with a faster website that returned the same results and did not show the work. We also found that people preferred websites that showed them the work over ones that did other things to distract from the wait—like providing entertaining pictures of their destination, promotional messages about other services offered by the website, or an interactive game of tic-tac-toe. None of those types of approaches made the service seem more valuable.

Why does operational transparency seem to have this unique power? We surveyed people who have (and have not) been given a glimpse behind the curtain in services as varied as restaurants, retail, and online dating to learn how operational transparency changes their perceptions. We found that when

people could see the work that was going on behind the scenes, they perceived that more effort went into the delivery of the service. They also believed that the service provider had more expertise and was being more thorough. They appreciated that effort and quality, and they in turn valued the service more.

In retail, for instance, Bhavya Mohan (of the University of San Francisco), Leslie John (of Harvard Business School), and I studied what happened when an online retailer added an infographic highlighting the costs and processes involved in manufacturing various products. For example, a wallet that sold for $115.00 included costs for raw materials ($14.68), construction ($38.56), duties ($4.26), and transportation ($1.00). Revealing the costs enabled the company to showcase to customers the otherwise hidden work that went into creating the wallet. In the process, of course, it also revealed that customers were paying $115.00 for something that cost $58.50 to make. The company further informed customers that its 1.9x markup compared favorably with the 6x markup charged by competing retailers—whose prices for similarly constructed items were higher. We found that sales of the wallets with operational transparency went up by 26% relative to wallets where the costs were not shared.

In subsequent experiments, we've learned that voluntarily providing operational transparency not only increases sales but also increases people's trust and satisfaction—even in settings where trust is otherwise low, such as government services.

According to the Pew Research Center, 73% of Americans in 1958 reported trusting government to do the right thing at least most of the time; today a paltry 20% do. So-called sunshine laws require a minimum level of transparency by elected officials and policy makers about certain of their activities, but those laws are not meant to spotlight the often invisible work that government does on a daily basis to create value in citizens' lives—such as disposing of trash, filling potholes, cleaning up graffiti, and fixing broken streetlights.

In 2009, Boston's local government developed a smartphone app called Citizens Connect (now BOS:311), which enables residents of the city to submit public service requests. Using the app, a resident can take a photo of a problem they want to report, such as a pothole, and the picture will automatically be geotagged using the phone's GPS and sent to the public works department. My colleagues Ethan Porter (of George Washington University), Michael Norton (of HBS), and I partnered with the City of Boston and Code for America in 2014 to study how showing the work being performed affected people's perceptions of government. We found that when people interacted with a website that showed images of the work being requested and performed, they became significantly more trusting and supportive of the government than if they interacted with a website that merely provided a tally of issues being reported and resolved. What's more, when the city took things a step further and asked employees to take

photos of the work they were doing and share them with the people who submitted the original requests, users became considerably more engaged, increasing the number of requests they made on a monthly basis by 60% and reporting issues in 40% more categories. Increased citizen engagement enabled Boston's government to allocate more workers to solving problems and fewer to finding them, so more work could get done.

The thoughtful application of transparency can create value even in settings where privacy is traditionally prized, such as health care. London Business School's Kamalini Ramdas and Nazli Sonmez and I collaborated with doctors at Aravind Eye Hospital, in Pondicherry, India, to study an application of operational transparency in delivering care to patients with glaucoma—an eye disease that is the second leading cause of blindness and afflicts some 12 million Indians. Some patients in our study were given appointments with their doctors in accordance with the hospital's normal protocol. Others were given shared appointments with three or four other patients. At the shared appointments, patients were able to see what the doctor could see when examining the eyes of others and hear the questions asked by other patients. Results from our ongoing collaboration suggest that patients who have shared medical appointments are more satisfied and engaged during their experience, are more likely to ask questions, learn more from the interactions, are more compliant with their

prescriptions, and are more likely to return for follow-up care than patients who have traditional one-on-one appointments with their doctor.

Although companies generally strive to make services appear as effortless as possible, examples of organizations beginning to experiment with various forms of operational transparency are becoming more abundant. When customers use an ATM to withdraw money from their BBVA bank accounts in Spain, the ATM's full-color screen displays visual representations of the currency being counted, sorted, and arranged for distribution. At most Starbucks drive-through locations in the United States, the intercom has been replaced with a video monitor and camera system. When customers place an order, they come face-to-face with the barista as he or she rings up the order and marks instructions on each cup. At Domino's, customers can use the company's Pizza Tracker app to watch as the kitchen workers prep, bake, and package the pizza for delivery.

NPR and the *New York Times* podcast *The Daily* are connecting listeners and readers with the otherwise obscure work involved in researching, producing, and delivering the headlines of the day. NPR posts live feeds from its studios, and *The Daily* interviews the paper's own reporters. In Detroit, the Mayor's Office has invested in the Neighborhood Improvement Tracker, a public-facing website that shows at a lot-by-lot level the many efforts being directed toward the city's recovery,

such as demolitions scheduled and completed to remove urban blight and building permits issued to enhance the community.

The evidence is clear: Operational transparency can fundamentally reshape the ways customers understand, perceive, and engage with the organizations that serve them. But what of employees?

Closing the Loop for Employees

Pioneering studies of service industries in the early 2000s found that a primary driver of satisfaction among employees is the knowledge that their company is delivering results to happy customers. Indeed, a 2007 study led by Adam Grant, an organizational psychologist and professor at Wharton, found that when call center agents soliciting donations for college scholarships actually met some of the students their work supported, their productivity and persistence skyrocketed. But what happens when the interaction between the customer and employee occurs in real time?

In 2012, Tami Kim (of the Darden School of Business), Chia-Jung Tsay (of University College London), and I ran an experiment in the Annenberg Hall dining facility at Harvard, which serves more than 3,000 meals every day. Annenberg was built in the late 1800s at a time when it was considered uncouth for diners to be able to see the work taking place in the kitchen. In that tradition, diners at Annenberg who desire

eggs, a fish sandwich, a hamburger, or some other grill item cooked their way must write their order on a piece of paper and hand it to an employee, who passes it through a small window into the kitchen, where a chef reads the order, cooks the item, and places it back in the window to be taken by an employee and given to the customer. The chefs can't see the customers, and the customers can't see the chefs.

We installed iPads with video-conferencing software—one at the order station, in view of the customers, and another in the kitchen, in view of the chefs. We then timed how long it took to make various dishes and measured both chef and diner satisfaction. When we turned on the iPads in a way that allowed only the chefs to see their customers, customer satisfaction with the food rose 14%. When we turned on the iPads so the customers could see the chefs too, satisfaction went up 22%, and the chefs worked 19% faster. One chef told us, "When [the customers] can see us [make their food], they appreciate it, and I appreciate that. It makes me want to improve."

Through surveys and additional experiments, we learned that when customers saw the chefs cooking their food, they perceived that more effort went into serving them, they appreciated the effort, and they valued the service more. When the chefs could see their customers—the people who were benefiting from their efforts—the work they were doing seemed more appreciated and impactful, making them more

satisfied with their jobs and more willing to exert effort. It was a virtuous cycle.

Consider another example: the Japanese train-cleaning company, Tessei, which I researched with Ethan Bernstein for an HBS case study. Tessei is charged with the Herculean task of cleaning the Shinkansen bullet trains during their brief stops at Tokyo station—1,000 seats in seven minutes. That's the equivalent of cleaning six Boeing 737s in less than half the time it typically takes to clean one. In the early 2000s, Tessei's employees were struggling to get the job done. Part of the challenge was that the work was underappreciated: Cleaning the bullet trains was known to be dirty and difficult, and so being a cleaner at Tessei was considered shameful in Japan. Accordingly, workers did whatever they could to escape the notice of customers. In 2005, a new leader, Teruo Yabe, revitalized the service, in part by promoting operational transparency among customers and employees. After the company changed employee uniforms from an invisible pale blue (which blended in with the body of the trains) to a vibrant red, passengers began to see and appreciate the work that these crews were doing, and after more interaction was instituted between the workers and the passengers, employees felt more appreciated and found a greater sense of purpose in their work. Employees began suggesting process improvements, and customers began chipping in to help tidy up their seats. There were quantifiable performance

improvements too; today a Tessei crew can clean a train in four minutes.

The India-based luxury hotel chain Oberoi Hotels takes operational transparency one step further, as I learned in my research for an HBS case study with Ananth Raman (of HBS) and Vidhya Muthuram (of the Blavatnik School of Government). Every employee in the company is preauthorized to spend up to Rs 1,500 (about US$25) to create moments of delight for guests. Whenever they learn of an opportunity to customize the service to improve a guest's experience, they're encouraged to act on it. The only stipulation is that employees must log what they have done so that the company and other employees can learn from their creativity. What has resulted is a feedback loop that fosters in employees a greater sense of purpose, helps customers feel better cared for, and improves organizational learning. Thanks in part to these efforts, Oberoi's properties routinely receive effusive reviews in customer surveys, and the company is perennially rated as one of the best luxury hotel brands in the world.

In contexts in which designing a face-to-face connection between employees and customers is impractical, technology can be used to successfully facilitate operational transparency. In 2013, Domino's piloted a feature called Domino's Live in one of its Salt Lake City locations, installing web cameras in the kitchen. Building on its Pizza Tracker app, customers ordering pizzas in Salt Lake could log on and watch a live feed of their

pizzas being made. As it turned out, tens of thousands of people from around the country logged on to watch other people's pizzas get made. Recognizing the potential, Domino's promoted Domino's Live on Facebook, and anytime someone clicked the "Like" button, a "Like Light" in the kitchen went on. This gave the pizza makers a signal that someone looking on appreciated the work they were doing. Although Domino's discontinued Domino's Live, the company added a feature to Pizza Tracker that enables customers to send notes of encouragement through the app to the people who are preparing their pizzas—prespecified messages such as "I don't know what I'd do without you" and "You are my pizza-making heroes." In a similar move, Uber recently updated its app to allow riders to close the loop with drivers—prompting them to send thank-you notes, along with tips, to the drivers after the ride is over. As one driver explained, "It makes my day to know when I've made somebody else's."

The Risk of Backfire

For all its benefits, operational transparency doesn't always deliver positive results. There are circumstances when it can repel customers and undermine employees. But even in such instances, managers should think twice before opting for complete opacity. Operational transparency should be carefully considered when:

It reveals things people genuinely don't want to see

Few may desire a behind-the-scenes look at trash collection or enjoying watching the dashcam footage of a violent police altercation. However, there's a difference between transparency that elicits the reaction "I'd rather not see that" and transparency that elicits the reaction "That should not happen." In the case of services that people aren't really interested in or find unappealing, companies should look for ways to use transparency to change the way people think about and engage with a service. For example, the city of Halifax, Nova Scotia, switched to clear trash bags in 2015 so that everyone could see what was being thrown away. Curbside waste collection fell by more than 30%, and recycling rates increased nearly 20%. When transparency causes people to object to what they see, organizations can draw on the experience to come up with alternative approaches that improve practice going forward. Dashcam footage of excessive violence by police departments has led to public outrage, but it has also improved oversight and accountability, sparked conversations that have led to policy change, and improved frontline training. "Out of sight, out of mind" may be more comfortable for everyone in the moment, but it rarely ensures the best long-term outcomes.

It engenders anxiety

Showing customers every step while their credit is being evaluated for a loan, or peering over employees' shoulders as they work, amplifies anxiety. Ethan Bernstein, of HBS, found that when curtains were put up around production lines at a Chinese cell phone manufacturer, productivity increased by 10% to 15%. Free from prying eyes, workers felt more focused and licensed to experiment with ways to improve standardized processes. What's more, workers felt safe to share ideas with one another, building team camaraderie and improving performance. When transparency makes us feel watched, it can hold us back; but when it helps us feel engaged, it can move us forward. For example, my HBS colleague Michelle Shell and I found that when customers who were transparently being evaluated for a loan were also provided with an easy way to contact a support person with questions throughout the process, the probability they would move forward with the loan, if offered, increased.

It shatters our faith in the relationship

When transparency reveals that a company isn't even-handed or that its practices violate implicit social norms, it makes customers understandably upset. Incidents of air rage—when an irate passenger causes a plane to land early—are higher on flights that have both a business class and an economy class and all passengers board from the front, forcing people in economy class to experience the disparity. This study,

conducted by Katherine DeCelles (of the Rotman School of Management) and Michael Norton, found that when the plane boards in the middle, so there's less transparency, the effect goes away. Or consider the ubiquitous marketing practice of personalizing ads. Tami Kim, along with Kate Barasz (of HBS) and Leslie John, found that when companies are transparent about targeting online ads on the basis of things we've revealed about ourselves, we appreciate the personalization. But when the transparency instead shows that they customize ads according to things they've *inferred* about us, it makes us upset. Customers also bristle when it's clear instead that companies are sharing their information with third parties without permission.

It destroys the magic

Sometimes we want to suspend our disbelief, and providing too much transparency would make that impossible. Retailers that sell high-end jewelry, musical instruments, or home decor often keep redundant inventory off the floor to give the pieces we see a special, one-of-a-kind mystique. The illusion that our ring or guitar or vase is unique enhances our experience. Likewise, even when it's 95 degrees outside, the cast member playing Mickey Mouse at Disneyland should keep the heavy, stuffy head of the costume on during the parade. Nothing can ruin the experience of make-believe like too much transparency. In other cases, we're fascinated to be in on the

secret. Factory tours and “how it’s made” shows are ubiquitous, and we clamor to watch bloopers and outtakes from our favorite movies. In fact, Disney offers a Backstage Magic experience for those who self-select into peeking behind the curtain.

It exposes an ineffective process

When transparency reveals employees who are incapable, indifferent, or powerless to deliver on the value proposition of the firm, customers can become incensed. Think back to the last service interaction you had where two employees were visibly chatting with each other instead of helping you. Or remember the last time your simmering frustration rose to a boil when a customer service rep repeated apologies for a problem over and over but had no means or authority to remedy the situation. Meanwhile, exposing employees to disenchanted and overtly negative customers, whom they have no hope of satisfying, can be a recipe for burnout. Agent turnover in many call centers, for example, exceeds 150% per year. Often situations like these arise when transparency hasn’t been designed to be reciprocal and to engender learning. Transparency that is accompanied by mechanisms to collect and learn from customer-provided feedback can accelerate, and create opportunities to celebrate, improvement.

It reveals that a company’s best efforts yield poor results

When people can see that a lot of behind-the-scenes effort went into creating an inadequate outcome, it reinforces their impression that the company is bad at what it does. In an experiment I conducted with Michael Norton, participants engaged with one of two online dating websites that gave them dissatisfying results. Participants perceived that the site that showed them how hard it was working was worse than the one that delivered the same bad result but didn't show the work. The impression was, "You tried so hard, and that's the best you could do? You must not be very good at your job." That said, when mistakes are made, timely transparency is still typically the best path. Customers may punish companies that fail to be transparent about missteps or errors, questioning the organization's motives for hiding the information. "Why did Equifax wait 40 days to inform 143 million people that their confidential information had been compromised?" customers might wonder. Or "Why did Facebook wait three years to disclose that Cambridge Analytica improperly accessed the records of 50 million users?"

It shows that the company's products or services are inferior to competitors'

A fundamental tenet of business still applies: If your customers find that your products are of poor quality, overly expensive, or otherwise less attractive than your competitors' offerings, they will do business elsewhere. Shwetha Mariadassou (of Stanford),

Yanchong Zheng (of MIT), and I found that such revelations are most damaging when a company's level of performance is seen as inferior to a competitor or industry benchmark. On the other hand, transparency that exposes a customer's own poor performance—for example, when your power company reports that you consume more electricity than your neighbors—can be a potent motivator of change. The effect can be especially powerful when the company reveals unflattering changes in your performance: You increased consumption by 5%, but your neighbors decreased consumption by an average of 3%.

It highlights a lack of progress

Uncertainty about our status makes our skin crawl. That's why progress bars are ubiquitous online, and why American, Delta, and United Airlines now update the status of people's bags throughout their journey, providing mobile alerts when bags have been scanned, loaded, off-loaded, placed in baggage claim, and so on. We like to have the feeling of moving forward, and transparency that demonstrates the opposite can be frustrating. For example, in a recent experiment, I found that when people who have been waiting for service can see that nobody has joined the queue behind them, they're significantly more likely to give up waiting than if they don't know whether anyone else has joined. Making visible their lack of progress from the end of the queue leaves them wondering whether continuing to wait is worthwhile. On the other hand,

when people who have been waiting for service are able to see that their time waiting has resulted in advancement from the end of the queue, they're significantly more likely to stay in line.

It reveals that the company's operations harm workers or the environment

News coverage of the 2013 collapse of Rana Plaza, which killed and injured thousands of Bangladeshi garment workers, and the 2010 Deepwater Horizon oil spill, which released millions of barrels of oil into the Gulf of Mexico, casts spotlights on inhumane working conditions and subpar environmental standards that reshaped corporate initiatives around supply chain sustainability. Visibility into such problems can cause a strong and swift customer backlash. To that end, transparency functions as a test of sorts: If you don't want people to see how you treat your employees or the planet, you probably need to make some changes. On the other hand, when transparency reveals that companies are operating sustainably, it can have a powerful effect.

Georgia Institute of Technology's Basak Kalkanci and I ran field experiments with Alta Gracia, an apparel manufacturer that pays a living wage to its workers in the Dominican Republic, and with Counter Culture Coffee, a North Carolina–based coffee roasting company that engages in environmentally sustainable practices. We collaborated with

the Looma Project to produce a short video showing footage of working conditions inside Alta Gracia's factory and featuring interviews with workers discussing the living wage that Alta Gracia pays. We produced a similar video highlighting Counter Culture Coffee's environmental sustainability practices, such as composting the chaff from its roasting process to reduce landfill waste. Showing these videos at point-of-sale kiosks increased the probability that customers would buy the company's products by roughly 20%, relative to merely showing brand videos.

It's deceptive

Transparency is helpful when it reveals work, but when the illusion of transparency is used to deceive or manipulate, it can backfire spectacularly. When customers call AT&T or Apple to request customer support, the companies' automated systems play the sound of typing between prompts to signal that work is being done. Customers understand these cues for what they are and do not mistake them for the sound of an actual person performing a task. However, companies can easily stray into dodgy territory. For example, several years ago, a company called Premier Health Plans used software to speak on behalf of telemarketing agents who had heavy accents. Calls would typically start off with the agent identifying "herself" as Samantha West and asking an initial question, prompting customers to think they were engaging with a live customer

service rep. However, awkward pauses between exchanges, the software's limited repertoire of phrases, and the mechanical word-for-word repetition that resulted during interactions caused skeptical customers to interrupt, asking, "Are you a robot?" Anticipating this possibility, the developers had included the recording of a disarming laugh and the response "I am a real person. Can you hear me OK?" Customers weren't buying it. Recordings soon emerged online of people interrogating Samantha West to expose her as a fraud.

Recently, Google announced its plans to roll out a much more sophisticated phone robot, called Google Duplex, that is fully automated and can pass as a human—calling restaurants and hair salons to make reservations and appointments on behalf of its users. The technology is breathtaking, and the potential for value creation is enormous, but unless Duplex is modified to be genuinely transparent, it's hard to imagine that those it deceives will be forgiving.

Bringing Operational Transparency to Your Organization

Given all the potential advantages and pitfalls of operational transparency, managers should be thoughtful about how they implement it. They should consider the following factors in designing their initiatives:

What to reveal?

A great place to start is to think about moments in the process that could be easily showcased with minimal effort. For example, one dessert-focused restaurant introduced operational transparency by suspending a tilted mirror from the ceiling above the pastry chefs who were plating and finishing desserts. Diners, whose views had been previously obscured by high counters and a bank of espresso machines, were captivated by their new window into the action.

Other opportunities for transparency can be found by considering what information already captured in the organization's databases would be appreciated by customers. For example, several years ago, as a part of its efforts to improve access to health care, the U.S. Department of Veterans Affairs began internally tracking how long veterans were waiting at each of its facilities to get an appointment to see a doctor. Recently, the agency made this information publicly available to patients on its website. Similarly, Quick Lane Tire and Auto Center, a nationwide auto repair company, has been experimenting with providing a digital information board in its waiting rooms that gives customers real-time updates about what's happening with their cars and the current status of the service queue.

When to reveal?

Transparency boosts value perceptions most when it reveals work as it is happening or just after it has been completed,

rather than showing work that has not yet occurred. In my research, I've found that customers are more satisfied when a travel site like Kayak shows its efforts to find a flight *as it searches* dozens of airlines than when it merely tells customers before they hit the "search" button *that it will search* dozens of airlines. In addition, consumers shouldn't be force-fed transparency. Rather, they should get to decide when they want to see more. For example, UPS receives 143 million package-tracking requests on a typical business day—which converts to an average of about seven lookups per package. These requests are made by customers who are actively curious about the status of particular packages and are tracking them at times of their choosing. Imagine if UPS instead called you at its own discretion seven times per shipment with a running progress report.

How to reveal?

Transparency implementations work best when they're visual—ideally giving customers actual windows into the process so that there's no question about the credibility of what's being shown. When this isn't possible, video or animated infographics and diagrams that provide a visual representation of the work boost the perception of value more than static imagery, which in turn outperforms text descriptions. Transparency also works best when it's voluntarily provided by companies; transparency that is wrung out of corporations as a

result of regulations, investor pressure, or other factors does not build trust.

Don't forget to close the loop. Transparency is the most beneficial when it's allowed to flow in both directions—from the customers into the operation and from the employees out to the customers. Forcing employees to toil in obscurity deprives them of seeing how their work is helping customers, reducing their feeling that their work is appreciated and undermining their motivation. What's more, transparency for employees can give them the information they need to customize service and help them learn better ways of operating.

In a sense, today's businesses have become victims of the global economy's immense productivity gains over the past two centuries. Consumers today rely on a dizzying array of products that are manufactured and distributed from all around the world and on services that are delivered with an intensifying frequency. But the apparently effortless abundance and convenience also make it easy for consumers to take work for granted and for employees to lose out on the learning and motivation that customer connections afford. With that in mind, businesses should stop reflexively hiding their operations for the sake of efficiency and instead

thoughtfully consider when and how to open them up to create more value for customers and employees alike.

Originally published in March–April 2019. Reprint R1902H

Disclosure: Ryan W. Buell has given paid lectures at Google and Uber in the past.

The Dual-Purpose Playbook

by Julie Battilana, Anne-Claire Pache, Metin Sengul, and Marissa Kimsey

CORPORATIONS ARE BEING PUSHED to change—to dial down their single-minded pursuit of financial gain and pay closer attention to their impact on employees, customers, communities, and the environment. Corporate social responsibility from the sidelines is no longer enough, and the pressure comes from various directions: rising and untenable levels of inequality, increasing evidence that the effects of climate change will be devastating, investors' realization that short-term profitability and long-term sustainability are sometimes in conflict. For reasons like these, a growing number of business leaders now understand that they must embrace both financial and social goals.

However, changing an organization's DNA is extraordinarily difficult. How can a company that has always focused on profit balance the two aims? It takes upending the existing business model. Not surprisingly, researchers have consistently found

that companies are quick to abandon social goals in the quest for profitability.

Yet some enterprises successfully pursue both. The U.S. outdoor-clothing company Patagonia, for example, which initially prioritized financial goals, has come to pursue social good more seriously over time. Others began with social goals but must earn revenue to survive. Grameen Bank, the Nobel Prize–winning microlender in Bangladesh, is an iconic example. We've spent a decade studying how socially driven businesses succeed, and what we've learned from in-depth qualitative studies and quantitative analyses may prove useful to traditional companies that want to adopt a dual purpose.

Our research reveals that successful dual-purpose companies have this in common: They take an approach we call *hybrid organizing*, which involves four levers: setting and monitoring social goals alongside financial ones; structuring the organization to support both socially and financially oriented activities; hiring and socializing employees to embrace both; and practicing dual-minded leadership. Taken together, these levers can help companies cultivate and maintain a hybrid culture while giving leaders the tools to productively manage conflicts between social and financial goals when they emerge, making the endeavor more likely to succeed.

Setting Goals, Monitoring Progress

Dual-purpose companies need to set goals along both financial and social dimensions and monitor performance on an ongoing basis.

Setting goals

Well-constructed goals are an essential management tool. They communicate what matters and can highlight what's working and what's not. These goals should go beyond mere aspirations to clarify a company's dual purpose for employees, customers, suppliers, investors, and regulators. Companies may need to experiment their way to a goal-setting model that works for them—something Grameen Veolia Water has managed by continually recalibrating its activities around explicit aims.

The company, which provides safe water in Bangladesh, started in 2008 as a joint venture between Grameen Bank and the water services provider Veolia. Veolia, which traditionally works under government contracts, recognized that no local authorities were responsible for providing drinking water to rural areas at that time. The partnership aimed to fill this gap. Its board set two goals for the new business at the outset: to provide safe, affordable drinking water to the inhabitants of the rural villages of Goalmari and Padua over the long term, and to sustain operations from sales without relying on grants.

Idea in Brief

The Problem

Corporations are being pushed to dial down their single-minded pursuit of financial gain and pay closer attention to their impact on society. But how can a company balance the two?

The Research

The authors have studied companies around the globe that pursue financial and social goals simultaneously. They find that the successful ones build a commitment to both economic and social value into their core organizational activities.

The Solution

Companies that want to do well and do good should focus on four key management practices: setting and monitoring dual goals; structuring the organization to support both goals; hiring and socializing employees to embrace them; and practicing dual-minded leadership.

These two goals came into conflict. When managers realized how difficult it would be to break even if they sold water only to poor rural households at a very low price, they designed a new revenue-generating activity: selling water in jars to schools and businesses in nearby urban areas. At this point it might have been tempting to focus attention and resources on the profitable new market segment at the expense of the original one. But leadership did not drift. The venture's clearly stated social goal reminded board members and managers that urban sales were meant to subsidize village sales. Ultimately the former amounted to half the company's revenues, helping Grameen Veolia Water pursue its social goal.

No single playbook exists for setting social goals. But our studies point to two rules of thumb. First, *do the research.* Often

leaders try to set goals without developing a deep understanding of the specific social needs they aim to address—or of how they may have contributed in the past to the buildup of problems. Just as they conduct market research to identify opportunities for profit, they should study those social needs. Their research should involve the intended beneficiaries along with other stakeholders and experts.

Prior to launching operations, Grameen Veolia Water conducted major research to understand water issues in Bangladesh, interviewing public officials and health and water experts along with community organizations. Managers discovered that some rural populations suffered not only from drinking surface water contaminated with bacteria (the researchers' initial assumption) but also from drinking water from wells built in the 1980s. Some well water, although clear and tasteless, was naturally contaminated by arsenic and was a major source of cancers in adults and cognitive impairment in children. This information led the business to focus its activity in Goalmari and Padua, which suffered from both sources of contamination. The company thus defined its goal as providing permanent access to clean water for everyone in those villages.

Second, *set goals that are explicit and enduring* (though they may have to be updated in light of a changing environment). Impact would be limited if the village residents consumed clean water for just a few years; to achieve a significant positive

change in their health, they would need access to clean water over decades.

Monitoring progress

Just as important as setting goals is identifying and adapting key performance indicators (KPIs) in order to measure the achievement of specific targets, be they financial or social. While we know how to measure sales, revenue growth, and return on assets, no widely accepted metrics currently exist for many social goals (although more progress has been made on measuring environmental impact). Nonetheless, it is possible to set both financial and social KPIs successfully. Our research has found that companies succeed by dedicating substantial time and effort to developing a manageable number of trackable metrics during the goal-setting process and revisiting them regularly to assess their continuing relevance and adequacy.

At Grameen Veolia Water, managers consulted with members of the rural communities they sought to serve and with academic experts before formalizing four KPIs: the company's self-financing ratio (its ability to fund planned investments from its own resources), the number of villagers with access to its services, the rate of rural penetration, and the rate of rural regular consumption (which captures both financial and social performance). The four numbers are updated monthly to monitor operations, and the board discusses them quarterly to guide strategic decision making.

A learning mindset is essential for developing and using KPIs. A willingness to experiment and change on the basis of experience, whether their own or others', helps businesses better understand social problems and how to address them. Dimagi's approach to setting social performance metrics exemplifies this mindset. Founded in 2002 and led by Jonathan Jackson, one of its cofounders, Dimagi provides software that NGOs and governments can use to develop mobile apps for frontline health-care workers in developing countries. At first Dimagi's primary social metric was the number of active users, which was meant to indicate how many people the technology positively affected. Jackson hoped to improve this metric, because it failed to distinguish between those who actually used the data to improve service delivery to patients and those who collected but did nothing with it.

The company formed a dedicated impact team to refine the social KPI. After exploration, the team created a metric —"worker activity months"—to measure the number of health care providers who were actually applying Dimagi's technology, and it implemented internal data systems to track the metric across all projects. But Jackson soon realized that this, too, was flawed, because the outcome was beyond Dimagi's control: How workers used the software depended more on the actions of Dimagi's clients—NGOs and governments—than on its own.

After reaching out to other social enterprises for advice, Jackson reverted to the number of active users as the

company's primary social barometer, yet combined it with a new entity—an impact review team—that focused on qualitative quarterly analyses and discussions about the impact of all projects. These reviews ensure that a team doesn't focus unduly on the quantifiable aspects of a project (revenue, costs, completion dates) but also explores the effectiveness of its service delivery and how that could be improved to better support frontline health-care workers. The team discusses indirect forms of impact as well, such as helping organizations assess their readiness for digitization.

Other successful businesses also complement KPIs with in-depth qualitative assessments of their social performance. For example, the Brazilian impact investing firm Vox Capital hired Jéssica Silva Rios, an executive dedicated to understanding and measuring its impact, and recently made her a full partner. Some companies also incorporate external social indicators developed by independent NGOs such as the Global Reporting Initiative, the Sustainability Accounting Standards Board, and B Lab. For example, Vox Capital monitors whether its rating from the Global Impact Investing Rating System is above average in comparison with other funds in developing markets and adjusts the fees it charges investors accordingly.

Structuring the Organization

It's virtually impossible to succeed on financial and social fronts over the long run if the company isn't designed to support both. Achieving an effective design requires that you think about which organizational activities create economic value and which create social value, how those activities relate to one another, and how you'll try to balance them.

Aligning activities and structure

Some activities create social and economic value at the same time. Others create predominantly one kind of value. For activities that create both kinds, an integrated organizational structure usually makes sense. Otherwise the activities are often best managed separately.

Revolution Foods, founded in 2006 by Kristin Richmond and Kirsten Tobey, provides nutritious lunches to low-income students in the United States. Richmond and Tobey created the company to serve a social purpose, having witnessed how poor food options hold kids back in underfunded schools. Every time they sell a healthful meal to a school, two things happen: They enhance a child's health, and they make money. Their core activity thus creates both kinds of value. As a result, they opted for an integrated structure, with a single manager in charge of operational efficiency, business growth, and the promotion of child well-being. Account managers often engage students in nutrition education (either directly or through community organizations), introducing them to new foods and

collecting their feedback on taste. The exposure to healthful foods enhances the long-term wellness of students and supports sales at the same time.

In contrast, the French company ENVIE learned over time that it needed to decouple the two kinds of activities. Launched in 1984, it had the goal of reintegrating long-term unemployed people into the job market by hiring them on two-year contracts to collect and repair used appliances for sale in secondhand shops. The company also provides support and training in how to repair appliances, how to look for a job, how to write a CV, and how to interview. The resale of appliances is what creates economic value. The training to enhance individuals' ability to find jobs outside ENVIE creates social value, but it doesn't make the company more profitable—in fact, it increases costs.

In the early years, staff members were asked to do two jobs: give beneficiaries technical guidance on how to repair or dismantle appliances (economic value) and provide them with social support (social value). However, it was difficult to find supervisors with both social and technical expertise. Even when they had both, the supervisors struggled to balance the two dimensions of their jobs. ENVIE's founders accordingly decided to set up separate organizational units, one for social support and one for repair, to be overseen by social workers and technical experts respectively. This increased the company's effectiveness in generating both kinds of value.

Creating spaces of negotiation

The rub is that tensions inevitably arise—particularly in differentiated structures. Left unattended, they can bring an organization to a halt. The Bolivian microlender Banco Solidario provides a cautionary example. In the 1990s constant resentment and fighting between bankers (concerned with fees and efficiency) and social workers (concerned with the affordability of loans and the livelihoods of microentrepreneurs) essentially froze the company. Loan officers quit left and right, the number of active borrowers plummeted, and the profit margin dropped. We've found that successful dual-purpose companies avoid such paralysis by supplementing traditional organizational structures with mechanisms for surfacing and working through tensions. These mechanisms don't make the tensions disappear—rather, they bring them into the open by letting employees actively discuss trade-offs between creating economic value and creating social value. Such deliberation provides a powerful safety valve and can speed up effective resolution.

Consider Vivractif, another French work-integration company. Founded in 1993, it hires and trains the long-term unemployed at recycling facilities. Those responsible for achieving one kind of goal or the other at the company often did not see eye to eye. While production supervisors managed workers to meet recycling targets, social workers were eager to take them away from the floor for mentorship and job-search

training. The company set up quarterly meetings between the two groups so that they could discuss each beneficiary's progress and bring up coordination issues. Joint work planning allowed both to share important deadlines (such as for commercial deliveries or social trainings) and to find joint solutions to scheduling conflicts. This improved productivity and furthered the company's social goals.

Spaces of negotiation can be successful in large companies as well. In one multinational cooperative bank headquartered in Europe, decision makers representing each of the local branches collectively make strategic decisions only after iterative debate, during which different groups of employees are responsible for championing either the social or the financial objectives of the organization. When individuals speak up about issues, their assigned roles prevent tensions from becoming personal.

Hiring and Socializing Employees

Embedding a dual-purpose focus in an organization's DNA requires a workforce with shared values, behaviors, and processes. Hiring and socialization are crucial to getting that right.

Hiring

Employees in a company that pursues dual goals tend to be successful when they understand and connect with both the

business and the social mission. We've seen companies mobilize such people by recruiting three types of profiles: hybrid, specialized, and "blank slate."

Hybrid individuals arrive equipped with training or experience in both business and social-value fields, such as environmental science, medicine, social work, and so forth. Such people are able to understand issues in both camps and can connect with employees and other stakeholders of either orientation.

Jean-François Connan is a good example. He was recruited in the late 1980s by Adecco, one of the largest temp work groups in the world, because he had training in industrial maintenance and human resources and experience as a teacher and a mentor for at-risk youth. The company hired him to help address a long-standing problem: A large number of its temp workers lacked strong qualifications. Connan played a leading role in building a dual-purpose subsidiary for Adecco that helps the long-term unemployed reenter the job market by hiring them for temp jobs. His background lets him interact seamlessly with Adecco leaders and corporate clients as well as with local partners (such as nonprofits dedicated to youth mentorship) and those whom they seek to serve. Now he is the company's head of responsibility and social innovation.

But hybrid employees aren't always available and may not always be the best fit. Dual-purpose corporations often hire *specialized* talent, which allows them to tap into deep expertise

and networks in each area. The main weakness of this approach is that it is more likely to result in conflict between groups, which may not understand each other's norms, vocabularies, and constraints—especially if the organization separates economic activities from social ones. As a result, tensions and turnover in these companies tend to be higher than in those with an integrated structure, producing a negative effect on the bottom line.

To mitigate this at Dimagi, Jackson explains the primacy of the organization's social purpose on his very first recruitment call with a technical expert (such as a software developer). After hiring, he creates opportunities for the expert to learn about the social business through formal talks, informal office interactions, and even face-to-face fieldwork in the underserved communities with which Dimagi works. Vox Capital, too, has hired managers with technical capabilities (such as fund management) and no experience in a social-mission-driven environment. Yet it systematically screens applicants for their ability to embrace and thus adapt to the company's hybrid culture.

When companies recruit *blank slate* individuals, who have experience in neither business nor the social sector, they put them in entry-level jobs and help them acquire dual values and skills. The Bolivian microcredit lender Los Andes S.A. Caja de Ahorro y Préstamo, founded in 1995, took this approach, hiring university graduates with hardly any professional experience

to become loan officers. The sense was that they would embrace a hybrid organizational culture more readily than experienced employees might. Of course, this approach has limitations. Taking inexperienced staffers into an organization may lower productivity. It also requires a considerable investment in training.

Although recruitment strategies obviously must be adapted to specific HR needs, we have observed that hybrid employees tend to be particularly well-suited for managerial and coordination positions; specialists can contribute useful expertise as middle managers in differentiated structures; and blank slates do best in entry-level jobs, where training won't be too challenging.

Socialization

Once people are on board, socializing them can be daunting. Every employee needs to understand, value, and become capable of contributing to both financial and social goals in some form.

Formal approaches to socialization may include companywide events such as annual general assemblies and retreats where dual goals and values are explained, discussed, assessed, and put into perspective. Dedicated trainings can remind employees—particularly those who specialize in just one sector—of the interconnectedness of revenue-generating and social-value-creating activities. Job-shadowing programs

and other forms of experiential training can also purposefully bring different groups together. At Vivractif social workers spend at least one day a year alongside recycling supervisors, and vice versa, so that each can learn and relearn about the company from the other perspective.

Another example comes from Oftalmología salauno, a Mexican company cofounded in 2011 by Javier Okhuysen and Carlos Orellana to provide high-quality, low-cost eye care to people who can't otherwise afford it. Although the pair saw economic goals and social goals as connected, they observed that some doctors focused only on patient care, and some managers considered only costs. So they formulated a set of core tenets and shared them at a daylong training for all employees, which clarified the interrelatedness of the company's financial and social aspects and gave employees a shared language for discussing tensions. Okhuysen and Orellana later instituted such sessions for new hires and continue to reinforce the training content in day-to-day interactions.

Spaces of negotiation can be valuable informal socialization opportunities, too. At Vox Capital a weekly time slot allows anyone to pose a question if he or she feels that the company's practices don't align with the organizational mission and values or is witnessing financial-social tradeoffs. Employees haven't shied away from tough topics. Some have asked whether its investment portfolio sufficiently emphasizes the social

missions of the businesses, while others have questioned whether the company's approach to raising capital is ethical.

Such conversations pushed cofounder Daniel Izzo to think critically about Vox's principles. "First I thought, It doesn't matter as long as [investors] don't have a say in what we do," he says. "But then someone asked, 'Would you take a drug lord as an investor?' Of course not. So there is a line. But where do we draw it? Do you take money from companies involved in corruption scandals in Brazil? Or from sons and daughters of top executives in those companies?"

Similarly, Bernardo Bonjean, who founded the Brazilian microfinance organization Avante in 2012, instituted a monthly breakfast where employees could come together and ask him questions. He also shares what's on his mind in letters to employees, discussing everything from the company's KPIs to his concerns about cash flow in the coming months. Okhuysen and Orellana put posters showing a matrix of Oftalmología salauno's four core tenets—commitment, service, reach, and value—in every meeting room. They can refer to these tenets when decision points arise, supporting a shared language among employees.

To encourage questions from employees, it's important to create an environment where people feel safe raising contentious issues. And when employees see changes in thinking and processes result from these discussions, they know that what they say is valued.

Events and conversations aren't the only ways to socialize employees. Promotion and compensation are also important. At the multinational cooperative bank mentioned above, being promoted to general director of a local branch requires excelling in business development, cost reduction, and profit making while also demonstrating a clear adherence to the company's social goals and a willingness to work collaboratively. One candidate for promotion commented, "I have seen many brilliant people fail because they did not embrace our values enough."

Vox Capital, like several other companies we studied, bases individual bonuses on both financial and social performance. Furthermore, Izzo is clear that he does not want the economic inequality that Vox is trying to redress in Brazil reproduced inside the company itself, so the maximum difference between employees' highest and lowest salaries and bonuses is capped at a multiple of 10. (In the United States in 2017 the average ratio of CEO-to-worker compensation was 312:1, according to the Economic Policy Institute.) Other companies, such as Revolution Foods, use shared ownership to motivate employees and increase their commitment to dual performance. Any full-time employee can become a shareholder through stock options. Richmond and Tobey believe that sharing ownership with employees, many of whom live in the low-income communities the company serves, is integral to their social mission.

Practicing Dual-Minded Leadership

Leaders must manage the tensions that inevitably crop up on the path to achieving dual goals. These tensions often involve competition for resources and divergent views about how to reach those goals. Leaders must affirm, embody, and protect both the financial and the social side and address tensions proactively.

Making decisions

Strategic decisions should embody dual goals. Whereas goals reflect aspirations, decisions provide real evidence of leaders' commitment to achieving specific aims. The experience of François-Ghislain Morillon and Sébastien Kopp is a good example.

Morillon and Kopp created Veja in 2004 to sell sneakers made under fair trade and environmentally friendly conditions in small cooperatives in Brazil. When they realized that advertising accounted for 70% of the cost of a typical major brand's sneakers, they made the bold decision not to advertise at all. That allowed them to sell sneakers at a price comparable to what their bigger competitors asked despite having production costs five to seven times as high. To make up for the absence of traditional advertising, the company formed strategic partnerships with high-end fashion brands such as agnès b. and Madewell and stores such as the Galeries

Lafayette to increase media exposure, grow sales, and become profitable.

At first Veja's clients—shoe retailers accustomed to the marketing of major sneaker brands—were skeptical. So Veja trained salespeople to educate them about the benefits of its product for people and the environment. Clients and the media now view the "zero ads" decision as evidence of the founders' commitment to their social goals, ultimately both giving the company social impact and making it profitable.

Morillon and Kopp also decided to temper the company's growth, despite increasing consumer demand in the United States. They refused to lower their fair trade and environmental standards to sell more shoes. Instead they decided to set production targets in keeping with the capacity of their fair trade partners while working closely with them to increase that capacity, ensuring a growth rate compatible with financial sustainability. That decision demonstrated, to employees in particular, the genuine commitment of Veja's leaders to their dual goals. In making bold decisions, the cofounders both emphasized the company's priorities and created the conditions for achieving them. They also showed that it's possible to avoid one of the most common pitfalls for dual-purpose companies: prioritizing profits over society when the pressure is on.

Profit allocation is another important area of strategic decision making. Dividends can be capped to ensure that

financial goals don't overshadow social ones. When founding Oftalmología salauno, Okhuysen and Orellana pledged to reinvest 100% of their profits for at least seven years, so the investors they selected—a social impact fund, the World Bank, and a private wealth-management fund—knew that no dividends would be paid during that time. Okhuysen explains: "Our investors ultimately expect both financial and social returns on their capital. But the alignment between us around reinvesting profits to improve and grow our network of eye-care clinics has helped ensure that financial goals do not take precedence over our social purpose."

Engaging the board

In successful hybrid companies, board members serve as guardians of the dual purpose. Thus they must collectively bring a combination of business and social expertise to the table. Diversity on the board is important for drawing the organization's attention to both social and financial goals, yet it increases the risk of conflict, because members with different perspectives are more likely to differ as to the best course of action. We have seen some companies experience near-paralyzing governance crises when socially and commercially minded board members with similar levels of influence strongly disagree.

Yet other companies have managed to avoid such crises because a chair or an executive director systematically bridged

gaps between the two groups. By fostering regular interactions and information sharing between them, such leaders enabled the groups to develop mutual understanding. Recall the subsidiary Jean-François Connan founded at Adecco. He invited representatives from prominent local nonprofits to join the board as minority shareholders, enabling the company to benefit from their social expertise, networks, and legitimacy and helping to protect the company's social mission. His hybrid experience put Connan in a good position to bridge the gap between the two groups of directors, fostering common ground by constantly reminding each of the importance of the other.

Some major roadblocks to dual-purpose organizing are outside a company's control. Chief among them is that the business ecosystem is still set up to prioritize the creation of shareholder wealth. The Global Reporting Initiative, the Sustainability Accounting Standards Board, and B Lab, among others, have taken steps to overcome some of these barriers. Each of them has created metrics for tracking companies' impact on the lives of employees and customers, the communities served, and the environment, providing organizations with benchmarks. What is at stake is ensuring that companies don't pick and choose areas of social focus on the basis of convenience.

Rating agencies are only one part of the ecosystem, however. Although more changes are under way—such as awarding legal

status to public benefit corporations in the United States, community interest companies in the United Kingdom, and *società benefit* in Italy—the regulations, educational standards, investment models, and norms that govern the production of economic value and social value are still mostly distinct from one another. As an increasing number of companies engage in hybrid organizing, the systems that support business also need to change.

But changing organizations and the ecosystem that surrounds them is difficult. Companies must fight the inertia of inherited ways of thinking and behaving. Trade-offs and tensions are inevitable, and success is more likely when leaders address them head-on. The four levers we have outlined are meant to help.

Originally published in March–April 2019. Reprint R1902K

How CEOs Manage Time

by Michael E. Porter and Nitin Nohria

IN THE LEXICON OF MANAGEMENT, the CEO is the epitome of leadership. Yet surprisingly little is known about this unique role. While CEOs are the ultimate power in their companies, they face challenges and constraints that few others recognize.

Running a large global company is an exceedingly complex job. The scope of the organization's managerial work is vast, encompassing functional agendas, business unit agendas, multiple organizational levels, and myriad external issues. It also involves a wide array of constituencies—shareholders, customers, employees, the board, the media, government, community organizations, and more. Unlike any other executive, the CEO has to engage with them all. On top of that, the CEO must be the internal and external face of the organization through good times and bad.

CEOs, of course, have a great deal of help and resources at their disposal. However, they, more than anyone else in the organization, confront an acute scarcity of one resource. That resource is *time.* There is never enough time to do everything

that a CEO is responsible for. Despite this, CEOs remain accountable for *all* the work of their organizations.

The way CEOs allocate their time and their presence—where they choose to personally participate—is crucial, not only to their own effectiveness but also to the performance of their companies. Where and how CEOs are involved determines what gets done and signals priorities for others. It also affects their legitimacy. A CEO who doesn't spend enough time with colleagues will seem insular and out of touch, whereas one who spends too much time in direct decision making will risk being seen as a micromanager and erode employees' initiative. A CEO's schedule (indeed, any leader's schedule), then, is a manifestation of *how* the leader leads and sends powerful messages to the rest of the organization.

A crucial missing link in understanding the time allocation of CEOs—and making it more effective—has been systematic data on what they actually do. Research on that has tended either to cover a small handful of CEOs, like the 1973 study in which Henry Mintzberg closely observed five chief executives (some of whom led nonprofits) for five days each, or to rely on large surveys that cover short periods (such as our HBS colleague Raffaella Sadun's 2017 study based on daily phone surveys with 1,114 CEOs from a wide variety of companies in six countries over one week).

Our study, which we launched in 2006, offers the first comprehensive and detailed examination of CEO time use in

large, complex companies over an extended period. To date, we have tracked the time allocation of 27 CEOs—two women and 25 men—for a full quarter (three months) each. Their companies, which are primarily public, had an average annual revenue of $13.1 billion during the study period. These leaders were all participants in the New CEO Workshop, an intensive program that every year brings newly appointed CEOs of large companies to Harvard Business School in two cohorts of 10 to 12 each. In total just over 300 CEOs have attended it.

In the study each CEO's executive assistant (EA) was trained to code the CEO's time in 15-minute increments, 24 hours a day and seven days a week, and to regularly verify that coding with the CEO. The resulting data set reveals where, how, and with whom the CEO spent his or her time and on what activities, topics, and tasks. Because it also covers what CEOs do outside of work, we have visibility into how CEOs balance work and personal life. In all, we collected and coded data on nearly 60,000 CEO hours.

Idea in Brief

The Problem

Managing the immense demands on their time is one of the biggest challenges CEOs face. Yet knowledge about how CEOs actually use time is almost nonexistent.

The Study

The authors tracked the activities of CEOs at 27 large companies 24/7 for 13 weeks and then held intensive debriefs with them. The resulting data set offers deep insights not just into time management but into the CEO's role itself.

The Findings

Leaders must learn to simultaneously manage seemingly contradictory dualities—integrating direct decision making with indirect levers like strategy and culture, balancing internal and external constituencies, proactively driving an agenda while responding to unfolding events, exercising leverage while being mindful of constraints, focusing on tangible decisions and the symbolic significance of every action, and combining formal power and legitimacy.

After CEOs completed the time-tracking phase, we shared their data with them, comparing it with anonymized data of the other CEOs we had studied up to that point. These intensive debriefings often included the CEOs' reflections on the pressures they faced in managing time, and on their mistakes and lessons learned. We also shared our accumulated data with the participants in each New CEO Workshop. In our discussions, CEOs routinely described managing time as one of their greatest challenges. The observations, questions, and personal approaches to allocating time they shared further enriched our understanding.

In this article we will do three things:

First, we'll provide a descriptive analysis of the data. How much time do CEOs spend at work versus on personal activities? How much do they spend in meetings versus thinking and reflecting alone? How much do they rely on email versus face-to-face conversation? Do they spend more time inside the company or outside, more with customers or investors? We'll answer those questions—and many more.

Second, we will offer prescriptions for how CEOs can manage their time more effectively across their many responsibilities. One of our most striking observations is that the way leaders allocate their time varies considerably. (See the exhibit "Looking Beyond the Averages.") Some of this variation reflects differences in their businesses and management practices. However, many time allocation decisions, such as participation in company rituals that offer limited return, reflect legacy norms and cultures, as well as a CEO's own habits. In our debriefings the CEOs all acknowledged that there were important areas where they could be using their time better. On the basis of these discussions and those with the hundreds of other CEOs in our workshops, we are convinced that every leader can improve his or her time management.

Finally, we will reflect on what our rich data reveals about the overall role of the CEO. A CEO has to simultaneously manage multiple dimensions of influence, which all contain dualities, or seeming contradictions, that effective CEOs must integrate. Understanding this broader view of the role is essential to success and also provides an important perspective for managing time well.

While our research focuses on the CEO role in large, complex companies, its findings have implications for all leaders (including executives of nonprofits) looking for ways to use their time and influence more effectively.

The Job Is All-Consuming

CEOs are always on, and there is always more to be done. The leaders in our study worked 9.7 hours per weekday, on average. They also conducted business on 79% of weekend days, putting in an average of 3.9 hours daily, and on 70% of vacation days, averaging 2.4 hours daily. As these figures show, the CEO's job is relentless.

About half (47%) of a CEO's work was done at company headquarters. The rest was conducted while visiting other company locations, meeting external constituencies, commuting, traveling, and at home. Altogether, the CEOs in our study worked an average of 62.5 hours a week.

Why such a grueling schedule? Because it is essential to the role. Every constituency associated with a company wants direct contact with the person at the top. As much as CEOs rely on delegation, they can't hand off everything. They have to spend at least some time with each constituency in order to provide direction, create alignment, win support, and gather the information needed to make good decisions. Travel is also an absolute must. You can't run a domestic company, let alone a global one, from headquarters alone. As a CEO, you have to be out and about.

Making time for personal well-being

Given that work could consume every hour of their lives, CEOs have to set limits so that they can preserve their health and their

relationships with family and friends. Most of the CEOs in our study recognized that. They slept, on average, 6.9 hours a night, and many had regular exercise regimens, which consumed about 9% of their nonwork hours (or about 45 minutes a day). To sustain the intensity of the job, CEOs need to train—just as elite athletes do. That means allocating time for health, fitness, and rest.

We paid special attention to the 25% of time—or roughly six hours a day—when CEOs were awake and not working. Typically, they spent about half those hours with their families, and most had learned to become very disciplined about this. Most also found at least some hours (2.1 a day, on average) for downtime, which included everything from watching television and reading for pleasure to hobbies like photography.

The CEO's job is mentally and physically demanding. Activities that preserve elements of normal life keep CEOs grounded and better able to engage with colleagues and workers—as opposed to distant, detached, and disconnected. CEOs also have to make time for their own professional renewal and development (which our data showed was often the biggest casualty of a packed schedule). And they must be careful, as our colleague Tom DeLong puts it, not to become "like race car drivers and treat home like a pit stop."

They Work Face-to-Face

The top job in a company involves primarily face-to-face interactions, which took up 61% of the work time of the CEOs we studied. Another 15% was spent on the phone or reading and replying to written correspondence. The final 24% was spent on electronic communications.

Face-to-face interaction is the best way for CEOs to exercise influence, learn what's really going on, and delegate to move forward the multiple agendas that must be advanced. It also allows CEOs to best support and coach the people they work closely with. How a CEO spends face-to-face time is viewed as a signal of what or who is important; people watch this more carefully than most CEOs recognize.

Avoiding the lure of email

In theory, email helps leaders cut down on face-to-face meetings and improve productivity. In reality, many find it ineffective and a dangerous time sink—but one they have trouble avoiding. Email interrupts work, extends the workday, intrudes on time for family and thinking, and is not conducive to thoughtful discussions. CEOs are endlessly copied on FYI emails. They feel pressure to respond because ignoring an email seems rude.

CEOs should recognize that the majority of emails cover issues that needn't involve them and often draw them into the operational weeds. Conversely, emails from the CEO can create a downward spiral of unnecessary communication and set the wrong norms, especially if the CEO sends them late at night, on weekends, or on holidays. It then becomes easy for everyone in

an organization to fall into the bad habit of overusing electronic communications.

That's why setting proper expectations and norms for what emails the CEO needs to receive—and when he or she will respond—is essential. Norms are necessary for the others in the organization as well, to prevent email from having a cascading effect on everyone, wasting precious hours and intruding on personal time. One way for the CEO to stay ahead of the digital avalanche is to have an adept EA filter messages and delegate many of them to others before the CEO even sees them. In the end, though, there is no substitute for being disciplined about resisting the siren call of electronic communications. This is a topic our CEOs were often animated about, and best practices in this area are still emerging.

Some CEOs in our study have begun to use videoconferencing as an alternative to face-to-face meetings, especially to cut down on travel for themselves and for team members who might otherwise have to come to see them. Although such efficiencies should surely be sought, CEOs must never forget that at its core their job is a face-to-face one.

They Are Agenda Driven

CEOs oversee a large number of organizational units and work streams and countless types of decisions. Our research finds that they should have an explicit personal agenda and that most do. A clear and effective agenda optimizes the CEO's limited time;

without one, demands from the loudest constituencies will take over, and the most important work won't get done.

A good agenda sets priorities for the CEO's personal involvement over the coming period. But it is not unidimensional; rather, it is a matrix including both broader areas for improvement and specific matters that need to be addressed, and it combines time-bound goals with more open-ended priorities.

In our study we asked each CEO to describe the agenda he or she was pursuing during the quarter being tracked and to highlight the hours devoted primarily to advancing it. Every executive provided an agenda. We found that the CEOs invested significant time—43%, on average—in activities that furthered their agendas. Some were far more disciplined about this than others: Time devoted to the core agenda varied widely, ranging from 14% to 80% of leaders' work hours. Most CEOs we talked with agreed that the more time they spent on their agendas, the better they felt about their use of time.

Overall, we found that an explicit agenda is one of the CEO's most important tools for making progress on multiple work streams simultaneously, addressing differences in the rate of progress across priorities, and using time effectively despite the need to respond personally to unforeseen events.

Advancing the agenda

Keeping time allocation aligned with CEOs' top priorities is so crucial that we suggest that every quarter CEOs make a point of

looking back at whether their schedule for the previous period adequately matched up with their personal agenda. They should also update the agenda to reflect current circumstances.

CEOs can benefit from making their personal agenda explicit to others. Their EAs and leadership teams both need to know and understand it so that they can stay aligned with it. (See the sidebar "Four Behaviors of Great Executive Assistants.") This understanding will help team members assume ownership of the goals and priorities of the work the CEO needs them to drive.

Four Behaviors of Great Executive Assistants

EXECUTIVE ASSISTANTS PLAY A VITAL ROLE in shielding CEOs from distractions and unnecessary activities and ensuring that leaders' limited time is used well. We often hear CEOs say that a highly skilled EA can dramatically increase their efficiency and effectiveness, and our research supports that view.

EAs often feel conflicting pressures, however, that can result in poor scheduling choices. For instance, although they may recognize that CEOs need time alone, our study shows that many EAs believe that a full CEO calendar signals that they're doing their job. They tend to book back-to-back appointments, limiting time for spontaneous communications or solitary reflection. In addition, while EAs recognize that protecting a CEO's time is one of their most important duties, some have a human reluctance to say no to people (especially colleagues in the organization). That allows unessential meetings to creep into the CEO's day. Conversely, other EAs take their traditional role as gatekeeper too far, maintaining such tight control over access that their bosses risk being seen as aloof or inaccessible.

Finding the right balance in managing the CEO's time requires judgment and emotional intelligence. It also requires strong communication skills, because an EA speaks for the CEO and can affect how a leader comes across. In our research we have identified four key behaviors that drive better performance:

1. **Understand the leader's agenda.** CEOs should have a written agenda detailing their top priorities (updated quarterly) and should spend much of their time on activities that advance the agenda. It's critical that the EA internalize this agenda and use it as a lens through which each meeting request is viewed. The CEO's responsibility is to ensure that the EA knows the agenda and the importance of keeping the schedule aligned with it.

2. **Include all the relevant players.** Managers at all levels tend to complain about having too many meetings. One solution is to try keeping meetings small and inviting only those whose attendance is essential. However, good CEOs delegate well, and to do so they need their direct reports and affected managers to be present. Otherwise, extra rounds of communication and follow-up will be needed after meetings. Good EAs avoid that problem by getting the right players in the room to begin with.

3. **Recognize the value of spontaneity.** Most CEOs are overbooked. They would benefit from more time to walk the hallways and initiate unplanned interactions. They also need room to react to events that can't be anticipated; leaving some open time in the leader's day will help EAs avoid frequently canceling and rescheduling appointments.

4. **Zealously protect personal and family time.** EAs should recognize that the long hours, travel, and stress of the CEO job can take a toll. Time with family and friends, regular exercise, and opportunities to recharge and reflect are crucial to effectiveness and avoiding burnout. EAs' daily scheduling choices play an important part in helping CEOs maintain the balance they need to succeed over the long haul.

Dealing with unfolding developments

A good portion of our CEOs' time (about 36%, on average) was spent in a reactive mode, handling unfolding issues, both internal and external. For many chief executives, it is not immediately clear when and how to address such issues or how much time to devote to them. Say that a member of the CEO's senior leadership team leaves a meeting looking upset. Should the CEO follow up with that person right away to make sure everything is OK? Should the CEO just wait and let the team member cool off? Sometimes emerging problems seem small at first but balloon into larger distractions if the CEO doesn't attend to them. In other instances a CEO's intervention makes an issue bigger than it might have been. It's essential for CEOs to figure out appropriate responses to these unfolding situations.

Every now and then, CEOs find themselves dealing with a sudden, full-blown crisis—a product or safety failure, a hostile activist's bid, a serious cyberattack, or even an external catastrophe such as a tsunami or a terrorist attack. Most of our CEOs (89%) spent some time on crises. Though on average it was small (1% of work time during the quarter we tracked), the total amount spent varied a great deal among the leaders in our study. Crises can create make-or-break moments in a CEO's leadership. In dealing with them, CEOs need to be highly visible and personally involved; the response to such events can't be delegated. Showing genuine concern for the people affected,

avoiding defensiveness, holding everyone together, and creating confidence that the organization will not only survive but emerge stronger are some of the things CEOs need to do during these times.

Limiting routine responsibilities

A surprisingly significant fraction (11%, on average) of our CEOs' work time was consumed by routine duties. Such activities varied considerably across CEOs, running the gamut from review meetings to board meetings, earnings calls, and investor days.

Operating reviews are a major component of a CEO's routine tasks. Their number, frequency, and length ranged widely across the leaders we studied, and our discussions suggested that some CEOs—especially those who had been COOs—overinvested in reviews that could be delegated to direct reports.

The ability of CEOs to control what we term "have-to-dos" was also quite variable. Have-to-dos include rituals such as giving welcome talks to new employees. These can play an important symbolic role and help reinforce the company's values and culture. By thoughtfully choosing which of these events to attend, CEOs can set the tone of their relationship with the organization. Yet a CEO must be disciplined about ensuring that feel-good activities don't collectively take up more time than he or she can afford.

Our discussions suggest that CEOs need to take a hard look at every activity that falls into the routine and have-to-do categories. They must ask whether it serves an important

purpose or is simply a company habit, something instituted by the predecessor, or a carryover from the CEO's previous role.

They Rely Heavily on Their Direct Reports

A CEO's direct reports are the company's most senior executives and include some of its most skilled managers. They span all the key elements of the business and offer CEOs the greatest opportunity for leverage. The leadership team, working together, can be the glue that helps the CEO integrate the company and get the work done.

In our study about half (46%) of a CEO's time with internal constituencies was spent with one or more direct reports, and 21% of it was spent only with direct reports. The total time spent with direct reports ranged from a low of 32% of time with internal constituencies to a high of 67%. When we explored that variation, we found that CEOs were more likely to spend time with their reports present when they had greater confidence in them.

We found that it's critical for each member of the leadership team to have the capabilities to excel and earn the CEO's full trust and support. Any weaknesses in this group significantly reduce the CEO's effectiveness, because dealing with work that reports should have handled, and cleaning up after them, eats up valuable time. In fact, when our CEOs gathered as a group across cohorts to see how things were going after they had been in office awhile, their number one regret was not setting high-

enough standards in selecting direct reports. Many CEOs told us this was because they focused too much on the present and not enough on the future when they first stepped into the role. Direct reports who could manage the status quo were often not the ones who could help the CEO take the company to a new level.

The more CEOs can delegate to their leadership team, the better they generally feel about their use of time. It eases the burden of needing to get personally engaged, following up, and asking others to report back. Since CEOs see their direct reports so frequently, it is also easy to stay in touch with how things are going with matters they are handling.

Staying connected to other managers

The CEOs in our study also spent considerable time (32% of their time with internal constituencies, on average) with a broader group of senior leaders, often called the top 100 (plus or minus). Many in this group report to the CEO's direct reports. We found that time with this next level of leadership was well spent. The top 100 are often the driving force for execution in the organization, and direct contact with the CEO can help align and motivate them. These leaders are also key to succession planning: Some will be candidates to replace the company's most senior executives. Given that the people at this level are often a generation younger, a few may eventually even be candidates to succeed the CEO. So getting to know them personally can be very useful.

Not surprisingly, the CEOs in our study spent less time with lower-level managers (14%, on average) and even less time with rank-and-file employees (about 6%, on average). However, our research suggests that effective CEOs need to be careful to maintain a human face in the organization. They must stay approachable and find ways to meaningfully engage with employees at all levels. This not only keeps them in touch with what is really going on in the company but helps them model and communicate organizational values throughout the workforce.

Direct human contact with the rank and file also grounds CEOs and helps them understand employees' reality. CEOs face a real risk of operating in a bubble and never seeing the actual world their workers face. Relationships with employees at multiple levels also build a CEO's legitimacy and trustworthiness in the eyes of employees, which is essential to motivating them and winning their support.

Knowing what is going on

Spending time with the rank and file, and with savvy external frontline constituencies, is also an indispensable way to gain reliable information on what is really going on in the company and in the industry. This is a major CEO challenge. Some CEOs get frontline contact by walking the hallways and factory floors, and using mechanisms like periodic lunches, unscheduled visits, and carefully designed field trips to customer and company sites. Others use group interactions, such as town halls, to foster

genuine and open conversations with a large cross section of employees (rather than present slide decks). Our data indicates that CEOs have varying success in carving out time for such steps, however.

They Manage Using Broad Integrating Mechanisms

CEOs must avoid trying to do too much themselves. It just isn't possible for them to make or even ratify most decisions directly. Instead, effective CEOs put in place well-designed structures and processes that help everyone else in the organization make good choices. These inform, support, enable, and integrate the work of others while building the organization's capabilities.

The most powerful integrating mechanisms include strategy (on which CEOs in our study spent an average of 21% of their work time), functional and business unit reviews (25% of their time), developing people and relationships (25% of their time), matching organizational structure and culture with the needs of the business (16% of their time), and mergers and acquisitions (4% of their time).

Harnessing strategy

The CEO's single most powerful lever is ensuring that every unit —and the company as a whole—has a clear, well-defined strategy. Strategy creates alignment among the many decisions within a business and across the organization. By spending time on strategy, a CEO provides direction for the company, helps

make its value proposition explicit, and defines how it will compete in the marketplace and differentiate itself from rivals. Strategy also provides clarity on what the company will *not* do. A compelling strategy—if well understood throughout the organization—is motivating and energizing. And without clarity on strategy, the CEO will be drawn into too many tactical decisions.

In large, complex firms, CEOs can almost never spend enough time on strategy—they must constantly be working to shape it, refine it, communicate it, reinforce it, and help people recognize when they may be drifting from it. CEOs must also ensure that the strategy is renewed from time to time and based on changes in the environment. Portfolio choices such as divestitures, mergers, and acquisitions are critical to strategy, and a CEO must be personally involved with them.

Aligning organizational structure and culture

To foster appropriate decisions across the company, the organization's structure needs to be aligned with its strategy. Otherwise, the CEO will be drawn into endless adjudication among units. It can also become a big drain on the CEO and others if the organization is constantly lurching from one structure to another.

Culture—which encompasses an organization's values, beliefs, and norms—is another key CEO lever for reinforcing strategy and influencing how the organization as a whole goes about doing its work. CEOs can shape a company's culture in many ways, from

the time they spend talking about it at various forums, to personally living the valued behaviors, to recognizing, rewarding, and celebrating those who exemplify the desired culture while taking corrective action with those who don't. It is the CEO's job to champion the organization's culture and constantly look for opportunities to strengthen it.

Designing, monitoring, and improving processes

CEOs must ensure that the company's strategy is being well executed. This will occur when the organization has rigorous processes through which work—such as marketing plans, pricing, product development, and strategy development itself—is done. Good processes bring together the best organizational knowledge and keep the CEO from continually having to override decisions.

Formal reviews are essential to monitoring whether the company is delivering the required process performance. Though these consume a quarter of a CEO's total work time, they allow CEOs to track progress, provide regular feedback, uphold high standards, and ensure timely course corrections. Reviews are also necessary to make sure that lessons learned are used to enhance the various processes through which work gets done.

However, excessive participation in reviews can get the CEO too involved in the company's operations and mired in unnecessary details. We talked a lot with the CEOs in our study about this problem. We have found, again and again, that many have a hard time shedding the COO or president roles they may have previously held. Some also forget that their senior team

should bear the primary responsibility for many reviews and keep the CEO informed on a regular basis.

When CEOs fail to delegate reviews to direct reports who can handle them, they erode the autonomy and accountability of their management teams. That doesn't help CEOs get the best out of others.

Developing people and relationships

Building the company's leadership pipeline is an important CEO function in its own right. We have found that CEOs must be personally committed to and be involved in improving the quality of the company's leaders. They cannot just leave this task to HR. Leadership choices are also pivotal in shaping the company's culture. Who gets hired, promoted, or fired signals what is truly valued by the CEO and the company.

CEOs need to get the most out of an organization's talent, and to do that, they must forge personal connections. Our CEOs spent another quarter of their total work time in meetings that focused on building relationships. When trust is mutual, delegation comes more naturally, agreement is easier to reach, and less monitoring and follow-up are necessary. Good relationships also make people more likely to give you the benefit of the doubt when you need it—and to tell you the truth, which is invaluable at the top.

The time CEOs spend building social capital through a network of personal relationships has many benefits and is time well spent.

They Are Always in Meetings

CEOs attend an endless stream of meetings, each of which can be totally different from the one before and the one that follows. Their sheer number and variety is a defining feature of the top job. On average, the leaders in our study had 37 meetings of assorted lengths in any given week and spent 72% of their total work time in meetings.

Looking beyond the averages

How much do CEOs' practices differ? We've ranked the variation in their uses of time from the lowest to the highest.

	Degree of variation (standard deviation/mean)	
Meeting time	0.14	Low
Face-to-face interactions	0.14	
Time with internal constituencies	0.14	
Total workweek obligations	0.14	
One-hour meetings	0.21	
Scheduled time	0.22	
One-on-one meetings	0.24	
CEO-initiated meetings	0.28	
Weekend days worked	0.31	
Core agenda time	0.36	Medium
Meetings per week	0.36	
Electronic communication	0.38	
Time with direct reports	0.39	
Functional and business-unit review time	0.41	
People and relationship time	0.44	
Strategy time	0.48	
Time on organizational structure and culture	0.54	
Spontaneous time	0.59	High
Have-to-do time	0.59	
Time with other outside commitments	0.59	
Two-hour-plus blocks of alone time	0.70	
Time with rank-and-file employees	0.71	
Exercise time	0.89	
Time with investors	0.95	
Time with customers	1.10	

Making meetings shorter and more effective

CEOs need to regularly review which meetings are truly needed and which can be delegated, and to let go of ones they were

accustomed to in previous roles.

They should also take a hard look at meeting length. In our study, meetings that lasted an hour accounted for 32% of a CEO's meetings, on average. Meetings that were longer accounted for 38%, and shorter meetings, 30%. We found that the length of meetings was often a matter of organizational or personal habit or both—a default length (like one hour) was the norm.

"Standard" meeting times should be revisited with an eye toward shortening them. Doing this can significantly enhance a CEO's efficiency. In our debriefs, CEOs confessed that one-hour meetings could often be cut to 30 or even 15 minutes. Another good way to streamline things is to reset meeting norms: Every meeting should have a clear agenda, and to minimize repetition, attendees should come prepared. Effective CEOs spread these meeting norms throughout the organization.

Some CEOs were worried that they might appear standoffish if someone asked for an hour and the CEO (or the EA) offered 30 minutes. But we have found that meeting length is worth confronting. "Whatever they ask for, cut it in half," said one CEO.

Another important meeting attribute is the number and composition of attendees. One-on-one meetings were the most common (accounting for 42% of CEOs' meetings, on average), followed by meetings with two to five participants (21%). Although every CEO had meetings involving large groups of 50 or more—like town halls, leadership off-sites, or all-company meetings—these were infrequent (5% of meetings).

The emphasis on one-on-one and small group meetings makes sense for enabling delegation and relationship building, and allows confidentiality. But leaders should also look for opportunities to bring the right people together. An essential part of the CEO's role is to align various internal and external constituencies around a common understanding of issues, decisions, and action agendas. Having the right people in the room is a powerful way to build that alignment and avoid the need for repetitive, time-consuming interactions to bring everyone along.

Allowing for accessibility and spontaneity

The vast majority of our CEOs' time (75%, on average) was scheduled in advance. The CEOs initiated more than half (51%) of their meetings themselves.

While controlling the nature and number of meetings is essential, we also found that CEOs need to regularly set aside time for more spontaneous interaction (which represented 25% of their work time in our study). This frees up space for same-day appointments initiated by others, for opportune conversations or meetings, and for responding to unfolding events.

The amount of time our CEOs allowed for spontaneous meetings varied considerably, ranging from 3% to 61%. In our debriefings, CEOs who discovered that they had left little room for spur-of-the-moment meetings were often surprised and quick to recognize the need for change.

Spontaneity and accessibility enhance a CEO's legitimacy. Leaders whose schedules are always booked up or whose EAs see themselves as gatekeepers and say no to too many people risk being viewed as imperious, self-important, or out of touch. EAs play a key role in finding the right balance here.

Carving out alone time

It's also vital for CEOs to schedule adequate uninterrupted time by themselves so that they can have space to reflect and prepare for meetings. In our study, CEOs spent 28% of their work time alone, on average—but again, that varied a great deal, from a low of 10% to a high of 48%. Unfortunately, too much of this alone time (59% of it) was fragmented into blocks of an hour or less; too little (18%) was in blocks of two hours or longer. CEOs need to cordon off meaningful amounts of alone time and avoid dissipating it by dealing with immediate matters, especially their in-boxes. This proved to be a common problem among the CEOs in our study, who readily acknowledged it.

Given that time in the office is easily eaten up, alone time outside the office is particularly beneficial. Long-distance travel out of contact with the office often provides critical thinking time, and many CEOs swear by it. To capitalize on it, CEOs should avoid traveling with an entourage.

They Juggle Many External Constituencies

While the CEOs we studied spent the majority of their time (70%, on average) dealing with internal constituencies, a good chunk (30%, on average) was spent with outsiders: 16% with business partners (such as customers, suppliers, bankers, investors, consultants, lawyers, PR firms, and other service providers), 5% with the company's board of directors, and 9% on other outside commitments (service on other boards, industry groups, dealing with the media and the government, and community and philanthropic activities).

External constituencies can be just as demanding as internal ones. Everyone wants to talk to the CEO, and dealing with external stakeholders is time-consuming. It often involves longer workdays and time away from headquarters and from home. There is a risk of drifting toward outside commitments less tied to company success.

Finding time for customers

Most of our CEOs were dismayed to discover how little time they spent with their customers—just 3%, on average. It surprised some even more to learn that this was less than the amount they spent with consultants. The scant time devoted to customers is partly a function of the huge scope of internal responsibilities: As an executive ascends from managing a line of business (which involves more-frequent customer contact) to the job of leading the entire company, it is natural for customer-facing time to decline.

Nonetheless, the CEOs in our study clearly felt that 3% was too low. Customers are a key source of independent information about the company's progress, industry trends, and competitors. In the B2B space, meeting with customers' CEOs is highly valuable, since peer conversations can be very candid. In B2C companies, there are also rich opportunities for customer contact. For retail CEOs, for example, store visits—especially unannounced ones—are an indispensable way to talk to regular customers, not just the company staff.

Some CEOs systematically schedule time with customers. The CEO of a financial services firm in our study, for instance, aims to meet face-to-face with one customer a day. A manufacturing CEO allocates two days a month to customer visits. Other CEOs try to build customer visits into their travel. A habit of some type seems to be the most reliable way to ensure enough customer time.

Limiting time with investors

On average, our CEOs spent only 3% of their total work time on investors. Most of them found this surprising; they tended to believe they spent more. But while more time is likely to be better when it comes to customers, the same is not true with investors. Too many meetings with investors can easily become a time sink and can draw the CEO into trying to manage the stock price rather than focusing on business fundamentals. Staying in touch with a few key buy-side investors, doing quarterly calls, and holding an annual investor day may be all a CEO needs to do

—unless, of course, the company is dealing with serious investor unrest or activism. By and large, the CEOs in our study seem to have discovered such focus over time, after getting caught up early in their tenures in too much investor relations.

Limiting unrelated outside commitments

There is a real risk that CEOs will get distracted by outside activities not directly connected to the business, where they are in high demand and which often involve worthy community and social issues. Such activities consumed an average of almost 2% of the work time of the CEOs in our study. While CEOs should give back to their communities and play the role of business statespeople, they should carefully restrict the hours they personally spend on such activities and on participating in business groups. Though the CEO's presence can be important, overseeing and managing such work does not require the CEO and can be delegated to direct reports, for whom it is motivational and provides professional development opportunities.

Finding time for directors

All our CEOs understood the importance of spending time with their boards. In our study, interacting with directors accounted for 5% of CEOs' total work time, or 41 hours a quarter, on average. But again we saw significant variation: One CEO spent six hours with directors; another spent 165.

A CEO must never forget that the board is his or her boss and that "managing up" is vital to success. However, that involves more than board meetings, committee meetings, and board retreats; CEOs must find time to build meaningful one-on-one relationships with individual directors. This is essential to take advantage of each board member's particular expertise and perspective. At board meetings, it's often not clear where each director is coming from, but that knowledge is crucial in crises and when dealing with controversial topics. CEOs also need to keep the directors well informed and engage with them between meetings through newsletters and updates. A common understanding and alignment with the board is important in periods of stress or market challenge.

Dimensions of the CEO's Role and Influence

The data on CEOs' time use reveals that the sheer complexity of their role—the myriad types of work, activities, and constituencies—is much greater than has previously been documented or perhaps even understood.

In examining the CEO's role, we have come to see that their work entails six dimensions of influence. Each involves a duality—a seeming contradiction, akin to yin and yang—that CEOs must manage simultaneously in order to be effective. (See the exhibit "Managing the Dimensions of CEO Influence.")

Managing the dimensions of CEO influence

Chief executives exert influence along six dimensions, each of which involves a duality, or seeming contradiction akin to yin and yang. Managing these dualities simultaneously is a hallmark of effective CEOs.

Direct	**Indirect**
The CEO is directly involved in numerous agendas and makes many decisions.	The CEO also exerts much influence over the work of others, using integrative mechanisms, processes, structures, and norms.
Internal	**External**
The CEO works with the senior team and with employees at all other levels to get all the organization’s work done.	The CEO also engages myriad external constituencies, serving as the face of the company, and must bring these external perspectives to the organization.
Proactive	**Reactive**
The CEO must articulate a sense of purpose, have a forward-looking vision, and lead the company to greater success.	The CEO must also respond to events as they unfold, from daily issues to full-blown crises that will prove to have a major impact on the company’s success.
Leverage	**Constraints**
CEOs’ position and control of resources give them immense clout.	CEOs are constrained by the need to build buy-in, bring others along, and send the right message.
Tangible	**Symbolic**
The CEO makes many decisions about concrete things like strategic direction, structure, resource allocation, and the selection of key people.	Much of CEOs’ influence proves to be intangible and symbolic; their actions set the tone, communicate norms, shape values, and provide meaning.
Power	**Legitimacy**

CEOs hold formal power and authority in the company that is reinforced by their competence and track record.

CEOs' influence also rests on legitimacy that comes from their character and the trust they earn from employees through their demonstrated values, fairness, and commitment to the organization.

First, CEOs clearly have *direct* influence over many issues and decisions, as their numerous reviews and one-on-one meetings reveal. However, the inherent limits on CEOs' time and knowledge mean that much of their influence must also be *indirect.* Good CEOs are very much in charge but work through others using strategy, culture, and effective organizational processes that drive sound analysis and alignment across the organization. CEOs need to learn how to marry direct and indirect influence.

Second, much of a CEO's work necessarily involves *internal* constituencies and managerial tasks, and our data verifies the overwhelming amount of such work to be done. However, CEOs are unique in the degree to which they must also engage and influence numerous *external* constituencies and represent the company to the world. Effective CEOs connect their internal and external roles by bringing outside perspectives into the work of the company. They also need to make sure outside constituencies understand the company's work and value.

Third, much of a CEO's work is inherently *proactive*: It involves anticipating problems, gathering the facts, conducting analyses,

and making sound and timely choices. Here, the CEO sets and drives the agenda. However, *reacting* well to unplanned and unforeseen events and crises is some of the most important work CEOs do. Choices here, and the CEO's personal presence or lack of presence, can have major consequences both outside and within the organization. Such periods can make or break a company and the CEO's own capacity to lead.

Fourth, while CEOs have a great deal of *leverage* to exert because of their position in the hierarchy and access to resources, they also face numerous—and often unrecognized—*constraints* and complexities in exercising that leverage. They are constrained in how often they can overturn decisions that have been brought to them for approval or how quickly they can drive changes without securing the support and buy-in of their senior team and board of directors. They must identify the group or people who are needed to bring about a change and then figure out how to win over the leader that will mobilize them. CEOs must find the right balance between taking full advantage of the leverage they possess, while being equally sensitive to the constraints they must navigate and the constituencies they must bring along. Otherwise, resistance will emerge and come back to bite them.

Fifth, while much of the CEO's influence is highly *tangible*, involving decisions about things like strategic priorities, budget targets, and people selection, some of the CEO's greatest influence is *symbolic*. This comes from the meaning people attach

to a CEO's actions. What CEOs do (and don't do), including everyday things like how they dress, what cars they drive, where they park, where they eat, and whom they talk to and how—always sends implicit messages to the company and its constituencies. Everything a CEO does affects what the organization focuses on, its norms of behavior, and its culture and values. The symbolic effects of CEOs' choices can reach even further than their specific actions.

What Do CEOs Actually Do?

WHILE WE REALIZE THAT CORPORATE leaders are really busy, we know surprisingly little about their day-to-day schedules. To fill that gap, in 2006 Harvard Business School professors Michael Porter and Nitin Nohria began asking participants of their New CEO Workshop to track their use of time, 24/7, for 13 weeks. The data on these pages, which were created with assistance from Harvard Business School research associate Sarah Higgins, summarizes the information gathered on how 27 CEOs spent a total of nearly 60,000 hours. Here is how they allocated their time, on average, among various activities, places, priorities, meetings, and constituencies.

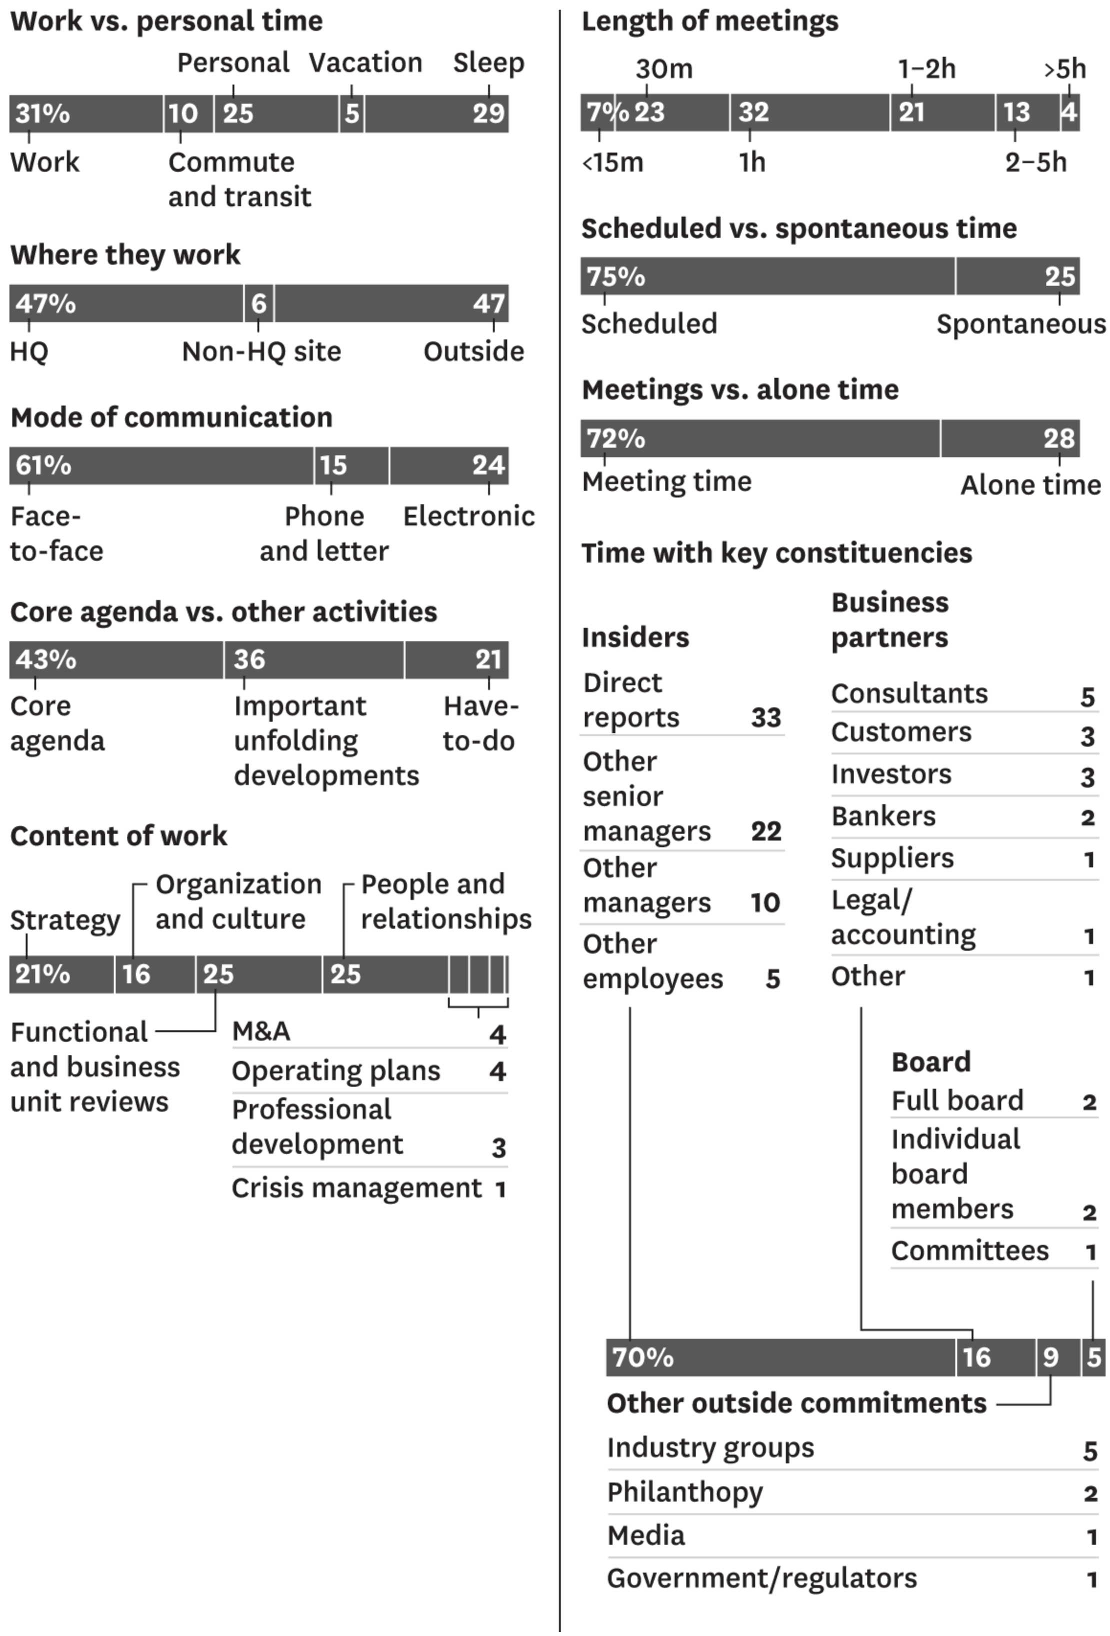
Work vs. personal time
Personal
Vacation
Sleep
31%
10
25
5
29
Work
Commute and transit
Where they work
47%
6
47
HQ
Non-HQ site
Outside
Mode of communication
61%
15
24
Face-to-face
Phone and letter
Electronic
Core agenda vs. other activities
43%
36
21
Core agenda
Important unfolding developments
Have-to-do
Content of work
Strategy
Organization and culture
People and relationships
21%
16
25
25
Functional and business unit reviews
M&A 4
Operating plans 4
Professional development 3
Crisis management 1
Length of meetings
30m
1–2h
>5h
7%
23
32
21
13
4
<15m
1h
2–5h
Scheduled vs. spontaneous time
75%
25
Scheduled
Spontaneous
Meetings vs. alone time
72%
28
Meeting time
Alone time
Time with key constituencies
Insiders
Direct reports 33
Other senior managers 22
Other managers 10
Other employees 5
Business partners
Consultants 5
Customers 3
Investors 3
Bankers 2
Suppliers 1
Legal/accounting 1
Other 1
Board
Full board 2
Individual board members 2
Committees 1
70%
16
9
5
Other outside commitments
Industry groups 5
Philanthopy 2
Media 1
Government/regulators 1

Sixth, CEOs hold a great deal of formal *power* and authority, and exercise it in the many ways we have described. However, power, authority, competence, and even results are insufficient to truly ensure their success. Effective CEOs combine formal power and authority with *legitimacy*. CEOs achieve legitimacy when employees believe in them as people and as leaders. They earn legitimacy in multiple ways—by demonstrating values, ethics, fairness, and a selfless commitment to the company and its people, among other things. Legitimacy gives rise to motivation that goes far beyond carrying out orders and can lead to extraordinary organizational performance. CEO time allocation, then, is not simply a matter of what happens in meetings and decision-making processes. It reflects the far broader set of ways in which the CEO as an individual engages with the organization and its people.

In managing across these six dimensions of influence, it is easy for CEOs to overlook the less direct, less top-down, less tangible, and more human aspects of their work. Without this awareness, though, CEOs give up some of their most powerful levers for driving change.

Why Good Leaders Matter

Countless concepts, tools, and metrics have been developed to help leaders manage well. However, our study of what the CEOs of large, complex organizations actually do—as manifest in how they spend their time—opens a new window into what

leadership is all about and into its many components and dimensions. Being the CEO is a highly challenging role, and it is difficult to do it well.

The success of CEOs has enormous consequences—good or bad—for employees, customers, communities, wealth creation, and the trajectory of economies and even societies. Being a CEO has gotten harder as the size and scope of the job continue to grow, organizational complexity rises, technology advances, competition increases, and CEO accountability intensifies. The ideas we have introduced here aim to provide current and future leaders, who must bear this enormous responsibility, with a broader understanding of their role and how to best use their most important resource: their time.

Originally published in July–August 2018. Reprint R1804B

When No One Retires

by Paul Irving

BEFORE OUR EYES, THE WORLD is undergoing a massive demographic transformation. In many countries, the population is getting old. Very old. Globally, the number of people age 60 and over is projected to double to more than 2 billion by 2050 and those 60 and over will outnumber children under the age of 5. In the United States, about 10,000 people turn 65 each day, and one in five Americans will be 65 or older by 2030. By 2035, Americans of retirement age will eclipse the number of people aged 18 and under for the first time in U.S. history.

The reasons for this age shift are many—medical advances that keep people healthier longer, dropping fertility rates, and so on—but the net result is the same: Populations around the world will look very different in the decades ahead.

Some in the public and private sector are already taking note—and sounding the alarm. In his first term as chairman of the U.S. Federal Reserve, with the Great Recession looming, Ben Bernanke remarked, “In the coming decades, many forces will shape our economy and our society, but in all likelihood no

single factor will have as pervasive an effect as the aging of our population." Back in 2010, Standard & Poor's predicted that the biggest influence on "the future of national economic health, public finances, and policymaking" will be "the irreversible rate at which the world's population is aging."

The world is getting older

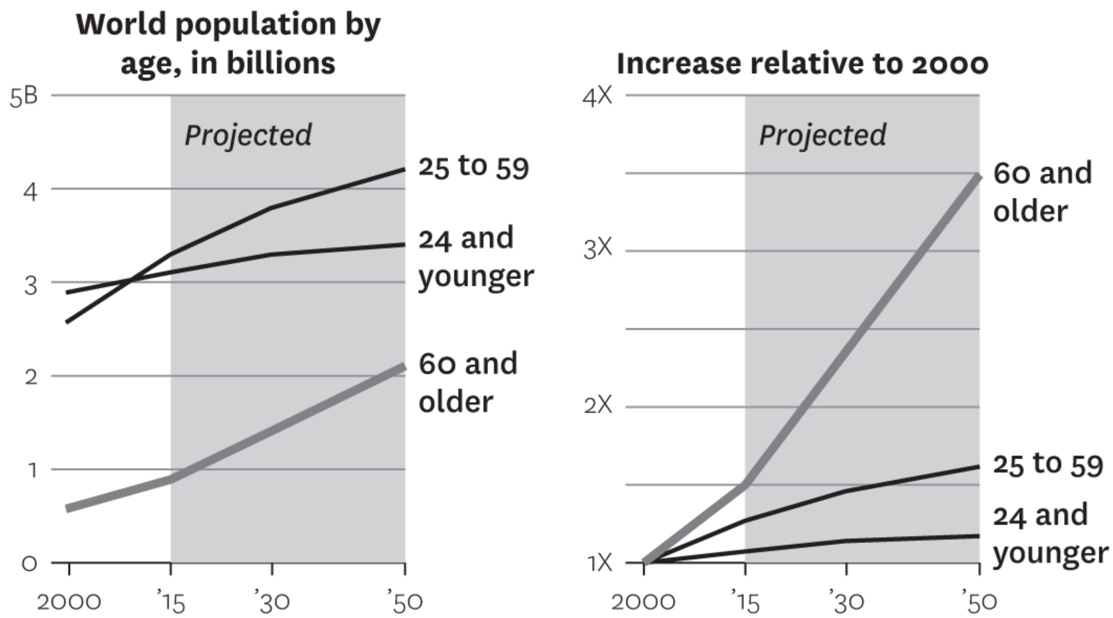

Source: United Nations, "World Population Prospects: The 2015 Revision."

This societal shift will undoubtedly change work, too: More and more Americans want to work longer—or have to, given that many aren't saving adequately for retirement. Soon, the workforce will include people from as many as five generations ranging in age from teenagers to 80-somethings.

Are companies prepared? The short answer is "no." Aging will affect every aspect of business operations—whether it's talent

recruitment, the structure of compensation and benefits, the development of products and services, how innovation is unlocked, how offices and factories are designed, and even how work is structured—but for some reason, the message just hasn't gotten through. In general, corporate leaders have yet to invest the time and resources necessary to fully grasp the unprecedented ways that aging will change the rules of the game.

What's more, those who *do* think about the impacts of an aging population typically see a looming crisis—not an opportunity. They fail to appreciate the potential that older adults present as workers and consumers. The reality, however, is that increasing longevity contributes to global economic growth. Today's older adults are generally healthier and more active than those of generations past, and they are changing the nature of retirement as they continue to learn, work, and contribute. In the workplace, they provide emotional stability, complex problem-solving skills, nuanced thinking, and institutional know-how. Their talents complement those of younger workers, and their guidance and support enhance performance and intergenerational collaboration. In encore careers, volunteering, and civic and social settings, their experience and problem-solving abilities contribute to society's well-being.

Idea in Brief

Before our eyes, the world is undergoing a massive demographic transformation. In many countries, the population is aging rapidly. In

the United States, about 10,000 people turn 65 each day, and one in five Americans will be 65 or older by 2030. This societal shift will affect every aspect of business operations, but corporate leaders have not yet grasped the unprecedented ways that an aging workforce will change the rules of the game. Those who *do* think about the impacts typically see a looming crisis—not an opportunity. This article helps companies develop a "longevity strategy" for fostering a vibrant multigenerational workforce.

In the public sector, policy makers are beginning to take action. Efforts are under way in the United States to reimagine communities to enhance "age friendliness," develop strategies to improve infrastructure, enhance wellness and disease prevention, and design new ways to invest for retirement as traditional income sources like pensions and defined benefit plans dry up. But such efforts are still early stage, and given the slow pace of governmental change they will likely take years to evolve.

The Big Idea: The Aging Workforce

"When No One Retires" is the lead article of HBR's **The Big Idea: The Aging Workforce.** Read the rest of the series at hbr.org/aging:

- "Rethinking Retraining," by Willy C. Shih, Howard Rudnick, and Colleen Tapen
- "Caring for Your Company's Caregivers," by Sarita Gupta and Ai-Jen Poo
- "Retirement-Proof Your Company," by Peter Berg and Matt Piszczek
- "Just How Old Are We Getting?" by Ramsey Khabbaz and Matt Perry

- “What Happens to Younger Workers When Older Workers Don’t Retire,” by Nicola Bianchi, Jin Li, and Michael Powell

---

Companies, by contrast, are uniquely positioned to change practices and attitudes *now.* Transformation won’t be easy, but companies that move past today’s preconceptions about older employees and respond and adapt to changing demographics will realize significant dividends, generating new possibilities for financial return and enhancing the lives of their employees and customers. I spent many years in executive management, corporate law, and board service. Based on this experience, along with research conducted with Arielle Burstein, Kevin Proff, and other members of our staff at the Milken Institute Center for the Future of Aging, I have developed a framework for building a “longevity strategy” that companies can use to create a vibrant multigenerational workforce. Broadly, a longevity strategy should include two key elements: internal-facing activities (hiring, retention, and mining the talents of workers of all ages) and external-facing ones (how your company positions itself and its products and services to customers and stakeholders). In this article, I’ll address the internal activities companies should be engaging in.

But first, let’s examine why leaders seem to be overlooking the opportunities of an aging population.

The Ageism Effect

There's broad consensus that the global population is changing and growing significantly older. There's also a prevailing opinion that the impacts on society will largely be negative. A Government Accountability Office report warns that older populations will bring slower growth, lower productivity, and increasing dependency on society. A report from the Congressional Budget Office projects that higher entitlement costs associated with an aging population will drive up expenses relative to revenues, increasing the federal deficit. The World Bank foresees fading potential in economies across the globe, warning in 2018 of "headwinds from ageing populations in both advanced and developing economies, expecting decreased labour supply and productivity growth." Such predictions serve to further entrench the belief that older workers are an expensive drag on society.

What's at the heart of this gloomy outlook? Economists often refer to what's known as the dependency ratio: the number of people not typically in the workforce—those younger than 15 and older than 65—in a population divided by the number of working-age people. This measure assumes that older adults are generally unproductive and can be expected to do little other than consume benefits in their later years. Serious concerns about the so-called silver tsunami are justified if this assumption is correct: The prospect of a massive population of sick, disengaged, lonely, needy, and cognitively impaired people is a dark one indeed.

The global aging phenomenon

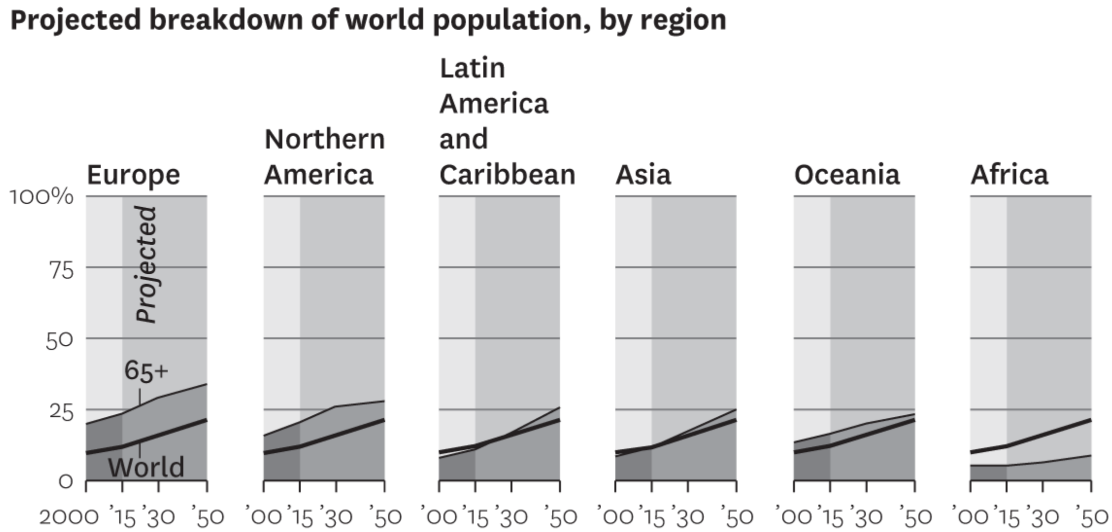

Note: Northern America consists of Canada and the United States.
Source: United Nations, "World Population Prospects: The 2015 Revision."

This picture, however, is simply not accurate. While some older adults do suffer from disabling physical and cognitive conditions or are otherwise unable to maintain an active lifestyle, far more are able and inclined to stay in the game longer, disproving assumptions about their prospects for work and productivity. The work of Laura Carstensen and her colleagues at the Stanford Center on Longevity shows that typical 60-something workers today are healthy, experienced, and more likely than younger colleagues to be satisfied with their jobs. They have a strong work ethic and loyalty to their employers. They are motivated, knowledgeable, adept at resolving social dilemmas, and care more about meaningful contributions and less about self-advancement. They are more likely than their younger

counterparts to build social cohesion and to share information and organizational values.

Yet the flawed perceptions persist, a byproduct of stubborn and pervasive ageism. Positive attributes of older workers are crowded out by negative stereotypes that infect work settings and devalue older adults in a youth-oriented culture. Older adults regularly find themselves on the losing end of hiring decisions, promotions, and even volunteer opportunities. Research from AARP found that approximately two-thirds of workers ages 45 to 74 said they have seen or experienced age discrimination in the workplace. Of those, a remarkable 92% said age discrimination is very, or somewhat, common. Research for the Federal Reserve Bank of San Francisco backs this up. A study involving 40,000 made-up résumés found compelling evidence that older applicants, especially women, suffer consistent age discrimination. A case in point is IBM, which is currently facing allegations of using improper practices to marginalize and terminate older workers.

There's more: Deloitte's 2018 Global Human Capital Trends study found that 20% of business and HR leaders surveyed viewed older workers as a competitive disadvantage and an impediment to the progress of younger workers. The report concludes that "there may be a significant hidden problem of age bias in the workforce today." It also warns that "left unaddressed, perceptions that a company's culture and

employment practices suffer from age bias could damage its brand and social capital."

The negative cultural overlay about aging is reinforced by media and advertising that often portray older adults in clichéd, patronizing ways. A classic example is Life Alert's ad from the 1980s for its medical alert necklace, immortalizing the phrase "I've fallen, and I can't get up!" Recent ads by E*TRADE and Postmates have also drawn criticism as ageist. A more subtle, but just as damaging example is the trumpeting of "anti-aging" benefits on beauty products as a marketing tool, suggesting that growing older is, by definition, a negative process.

Some companies are pushing back: In a recent video, T-Mobile's John Legere took on the topic of ageist stereotypes while promoting a T-Mobile service for adults age 55-plus. He chided competitors for what he called their belittling treatment of older adults in marketing campaigns that emphasize large-size phone buttons and imply that boomers are tech idiots. "Degrading at the highest level," Legere calls it. "The carriers assume boomers are a bunch of old people stuck in the past who can't figure out how the internet works. News flash, carriers: Boomers invented the internet."

Yet for the most part, employers continue to invest far more in young employees and generally do not train workers over 50. In fact, many companies would rather not think about the existence of older workers all. "Today it is socially unacceptable to ignore, ridicule, or stereotype someone based on their gender,

race, or sexual orientation," points out Jo Ann Jenkins, the CEO of AARP. "So why is it still acceptable to do this to people based on their age?"

Over the past decades, companies have recognized the economic and social benefits of women, people of color, and LGBT individuals in the workforce. These priority initiatives must be continued—obviously, we're not even close to achieving genuine equality in the corporate world; at the same time, the inclusion of older adults in the business diversity matrix is long overdue. Patricia Milligan, senior partner and global leader for Mercer's Multinational Client Group, observes, "At the most respected multinational companies, the single class not represented from a diversity and inclusion perspective is older workers. LGBT, racial and ethnic diversity, women, people with physical disabilities, veterans—you can find an affinity group in a corporation for everything, except an older worker."

The U.S. labor force is getting older, too

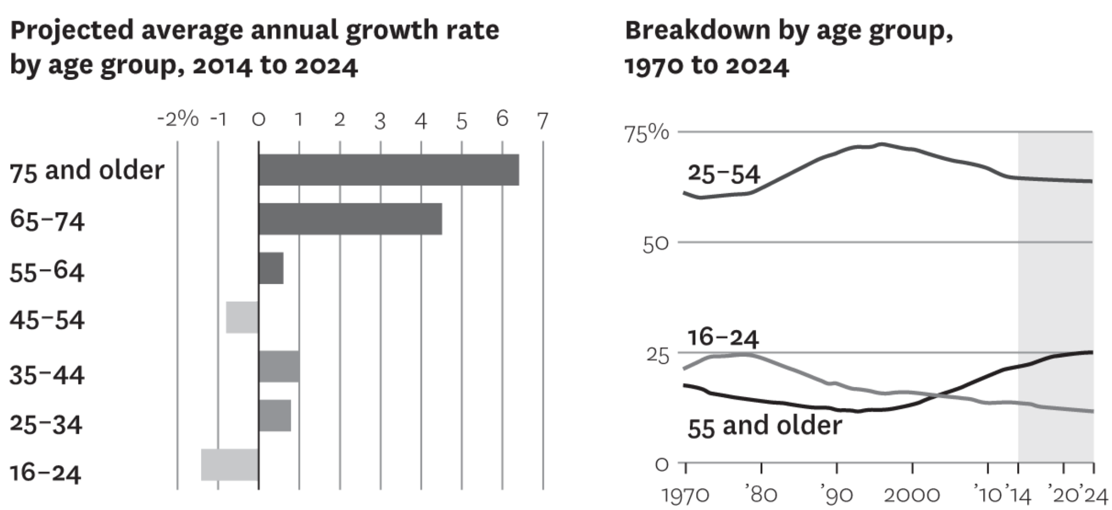

Source: U.S. Bureau of Labor Statistics.

Managing a Multigenerational Workforce

How can companies push past stereotypes and other organizational impediments to tap into a thriving and talented population of older workers? A range of best practices have been emerging, and some companies are making real progress. Each points to specific changes companies should be considering as they develop their own strategies.

Redefine the workweek

To start, you need to reconsider the out-of-date idea that all employees work Monday through Friday, from 9 to 5, in the same office. The notion that everyone retires completely by age 65 should also be jettisoned. Companies instead should invest in opportunities for creative mentorship, part-time work, flex-hour schedules, and sabbatical programs geared to the abilities and inclinations of older workers. Programs that offer preretirement and career transition support, coaching, counseling, and encore career pathways can also make employees more engaged and productive. Many older workers say they are ready to exchange high salaries for flexible schedules and phased retirements. Some companies have already embraced nontraditional work programs for employees, creating a new kind of environment for success. The CVS "Snowbird" program, for example, allows older employees to travel and work seasonally in different CVS

pharmacy regions. Home Depot recruits and hires thousands of retired construction workers, making the most of their expertise on the sales floor. The National Institutes of Health, half of whose workforce is over 50, actively recruits at 50-plus job fairs and offers benefits such as flexible schedules, telecommuting, and exercise classes. Steelcase offers workers a phased retirement program with reduced hours. Michelin has rehired retirees to oversee projects, foster community relations, and facilitate employee mentoring. Brooks Brothers consults with older workers on equipment and process design, and restructures assignments to offer enhanced flexibility for its aging workforce.

Reimagine the workplace

Your company should also be prepared to adjust workspaces to improve ergonomics and make environments more age-friendly for older employees. No one should be distracted from their tasks by pain that can be prevented or eased, and even small changes can improve health, safety, and productivity. Xerox, for example, has an ergonomic training program aimed at reducing musculoskeletal disorders in its aging workforce. BMW and Nissan have implemented changes to their manufacturing lines to accommodate older workers, ranging from barbershop-style chairs and better-designed tools to “cobot” (collaborative robot) partners that manage complicated tasks and lift heavier objects. The good news is that programs that improve the lives of older workers can be equally valuable for younger counterparts.

Mind the mix

Lastly, you need to consider and monitor the age mixes in your departments and teams. Many companies will need to manage as many as five generations of workers in the near future, if they aren't already. Some pernicious biases can make this difficult. For example, research shows that every generation wants meaningful work—but that each believes everyone else is just in it for the money. Companies should emphasize workers' shared value. "Companies pursuing Millennial-specific employee engagement strategies are wasting time, focus, and money," Bruce Pfau, the former vice chair of human resources at KPMG, argues. "They would be far better served to focus on factors that lead all employees to join, stay, and perform at their best."

By tapping ways that workers of different generations can augment and learn from each other, companies set themselves up for success over the long term. Young workers can benefit from the mentorship of older colleagues, and a promising workforce resource lies in intergenerational collaboration, combining the energy and speed of youth with the wisdom and experience of age.

PNC Financial Group uses multigenerational teams to help the company compete more effectively in the financial markets through a better understanding of the target audience for products. Pharma giant Pfizer has experimented with a "senior intern" program to reap the benefits of multigenerational collaboration. In the tech world, Airbnb recruited former hotel mogul Chip Conley to provide experienced management

perspective to his younger colleagues. Pairing younger and older workers in all phases of product and service innovation and design can create opportunity for professional growth. And facilitating intergenerational relationships, mentoring, training, and teaming mitigate isolation and help break down walls.

To begin this process, start talking to your employees of all ages. And get them to talk with each other about their goals, interests, needs, and worries. Young and old workers share similar anxieties and hopes about work—and also have differences that need to be better understood companywide. Look for opportunities for engagement between generations and places where older and younger workers can support one another through skill development and mentorship. After all, if everyone needs and wants to work, we're going to have to learn to work together.

To be clear, all of these changes—from flexible hours to team makeup—will require a recalibration of company processes, some of which are deeply ingrained. Leaders must ask, do our current health insurance, sick leave, caregiving, and vacation policies accommodate people who work reduced hours? Do our employee performance-measurement systems appropriately recognize and reward the strengths of older workers? Currently, most companies focus on individual achievement as opposed to team success. This may inadvertently punish older employees who offer other types of value—like mentorship, forging deep relationships with clients and colleagues, and conflict resolution

—that are not as easily captured using traditional assessment tools. Here, too, initiatives aimed at older workers can benefit other workers as well. For instance, research suggests that evaluating team performance also tends to boost the careers of employees from low-income backgrounds.

Turning a Crisis into an Opportunity

I'm admittedly bullish about the positive aspects of working longer and believe that company leaders can harness the opportunity of an aging population to gain competitive advantage. But I'm not oblivious to the challenges a longevity strategy poses. We're talking about initiating a massive culture change for firms—a change that must come from the top.

But ignoring the realities of the demographic shift under way is no longer an option. CEOs and senior executives will need to put the issue front and center with HR leaders, product developers, marketing managers, investors, and many other stakeholders who may not have it on their radar screens. This will take guts and persistence: Leaders must bravely say, "We reject the assumption that people become less tech-savvy as they get older" and "We will fight the impulse to put only our youngest employees on new initiatives." To genuinely make headway on this long-range issue, companies will have to make tough, and sometimes unpopular, decisions, especially in a world where short-term results and demands dominate leaders' agendas. But isn't that what great leaders do?

The business community has a chance to spearhead a broad movement to change culture, create opportunity, and drive growth. In doing so, companies will improve not only mature lives, but lives of all ages, and the prospects of workers for generations to come. This transformative movement to realize the potential of the 21st century's changing demography is the next big test for corporate leadership.

The Longevity Opportunity

by Paul Irving

AS THE GLOBAL POPULATION AGES, new consumer opportunities and markets will emerge. Every company should have a strategy for tapping into the needs, wants, and buying power of older customers.

In my previous article, I outlined the beginnings of an internal, employee-focused "longevity strategy" that organizations can use to reap the benefits of an aging workforce. In this one, I'll discuss its external, consumer-facing complement.

The market for products and services for older adults is already strong, and it will become even stronger. With distinct consumption habits and service needs, Americans over 50

accounted for $7.6 trillion in direct consumer spending and related economic activity in 2015, and controlled more than 80% of household wealth, according to a 2016 joint report from Oxford Economics and AARP. Further, a 2010 AARP survey reveals that 90% of older adults say they want to be able to remain in their own homes as they age. Envisioning how communities will respond to the needs of aging people will keep many more of them in their homes and contributing to the economy. Bank of America Merrill Lynch projects that the global spending power of those age 60 and over will reach $15 trillion annually by 2020.

But the potential for forward-thinking companies goes beyond just an interesting business opportunity. Older adults are poised to shape consumer and capital markets in the years ahead. The McKinsey Global Institute concludes that the 60-plus population, one of the few engines of global economic growth, is on track to generate half of all urban consumption growth between 2015 and 2030. "The Longevity Economy is redrawing economic lines, changing the face of the workforce, advancing technology and innovations, and busting perceptions of what it means to age," states the Oxford Economics/AARP report.

Across industries, there are multiple avenues for offering products and services that make a difference in people's lives. In the health sector, "gray is the new black," a Reuters piece observed. New offerings in biotechnology, devices, pharmaceuticals, and care services all target older consumers.

Research reveals that older adults dominate spending in 119 of 123 consumer packaged goods categories, spend more in grocery stores and purchase more new cars than any other age group, and account for 80% of luxury travel. The demographic is eager to spend on transportation, entertainment, food, and alcohol, representing an immense target market for fresh ideas and innovations.

The financial services industry has always catered to older people—primarily those planning for retirement. As customers and clients prepare for longer lives, retirement remains a powerful growth driver. But a financial market for older workers and entrepreneurs is rapidly expanding too. Driven by financial assets controlled by older investors, this segment of the longevity market simply cannot be ignored. People age 60 and over hold the majority of wealth worldwide and 70% of the disposable income in the U.S.

For most companies, longevity marketing is still in its early days. That must change. You need a strategy for older consumers, and identifying new opportunities and markets is just the first step. You and your employees also need to reconsider what you "know" about this population to avoid ageist and outdated messaging. Incorporate older employees into product planning, design, and communications to benefit from their experience and understanding. Use focus groups that include older participants to test products and services before they make it to market.

While most companies are in the early stages of developing their strategies, it's worth exploring what some industry leaders are up to.

Philips and Nestlé have fundamentally shifted their businesses to capitalize on trends that are driven in many ways by the aging population. Both global companies have focused their futures on health and wellness, recognizing the massive opportunities ahead. In partnership with the Global Social Enterprise Initiative at Georgetown University's McDonough School of Business, Philips is developing new technologies that can meet the needs of its older customers, including connected care solutions, safety applications, and cognitive health innovations. Patients and their doctors will be able to see, monitor, and share vital health information through secure devices, for example. Nestlé is investing in personalized diet and nutrition initiatives, and is broadening its portfolio by buying or acquiring stakes in health supplement and pharmaceutical companies.

Best Buy, with its recent acquisition of GreatCall, the provider of connected health and personal emergency services to the aging population, is focused on building relationships with older consumers. By gaining access to GreatCall's customer base, Best Buy can further penetrate the health services and monitoring business and grow through its supply chain efficiencies and marketing reach. Several analysts hailed the acquisition, praising Best Buy for recognizing the market's size and potential and the opportunity for a services business line to diversify the

company's offerings and counterbalance the margin pressures on electronics products.

Bank of America Merrill Lynch is training its customer-facing workforce to understand the needs of its aging clients. The bank recognizes that increasing longevity leads to new health care choices, housing issues, and questions about retirement and financial security. In partnership with the USC Leonard Davis School of Gerontology (which I'm affiliated with), the company's longevity training program teaches financial advisers about older people's experiences, priorities, and goals.

Uber and Lyft have developed programs to provide rides for older adults through age-friendly web tools, apps, and phone systems. Recognizing the importance of mobility to health and well-being, both ride-hailing companies are creating partnerships and scaling up efforts to facilitate access and ease of use for older adults and their families and caregivers. For example, both companies are working with call services (Lyft with GreatCall and Uber with RideWith24) to make it easier for older adults to book rides.

Intel is working on internet-of-things software that flags health concerns, and projects such as enabling wearables to analyze and communicate health data faster than ever through 5G internet connections. And Nest has begun modifying its line of smart home products to help older adults continue to live independently.

These are just a few examples of the longevity market's prospects and possibilities.

Finally, as longevity strategies are developed and implemented, companies must consider not only how their products and services are designed but also how they are promoted. The power of media and advertising should be used to reflect realistic images of older adults instead of stereotypes. Older consumers do not want to be patronized, but they do want their needs acknowledged, and companies can do this while emphasizing both positive and real aspects of aging. While many companies have a long way to go, some are getting it right. On the positive side, Unilever's Dove successfully employed its Pro-Age campaign, realizing significant market share increases. The smart marketing for Dannon's Activia yogurt is focused on the common issue of digestive health. And in 2017 *Allure* Magazine showed leadership when it announced that it would no longer use the term "anti-aging" to describe skin care or makeup.

We're still in the early stages of understanding what older consumers' needs are and how to address them. The aging population is diverse, and the answers are not simple; one size certainly does not fit all. But we do know that there's already a clear demand for products and services that can help people live longer, more-comfortable, and more-meaningful lives—and that are promoted without stigma or stereotype. This demand will grow rapidly in the coming decades, and companies that start meeting it now can reap a sizable dividend. It's a huge

opportunity, one that will have benefits both to their bottom lines and to society.

Originally published on hbr.org in November 2018. Reprint BG1806

About the Contributors

JULIE BATTILANA is the Joseph C. Wilson Professor of Business Administration at Harvard Business School and the Alan L. Gleitsman Professor of Social Innovation at Harvard Kennedy School.

ADAM BRANDENBURGER is the J.P. Valles Professor at the Stern School of Business, Distinguished Professor at the Tandon School of Engineering, and faculty director of the Shanghai Program on Creativity + Innovation at New York University.

ALISON WOOD BROOKS is an assistant professor at Harvard Business School.

RYAN W. BUELL is the UPS Foundation Associate Professor of Service Management at Harvard Business School.

PAUL R. DAUGHERTY is Accenture's chief technology and innovation officer. He is a coauthor of *Human + Machine: Reimagining Work in the Age of AI* (Harvard Business Review Press, 2018).

ROBIN J. ELY is the Diane Doerge Wilson Professor of Business Administration at Harvard Business School and the faculty chair of the HBS Gender Initiative.

JOSHUA GANS is the Jeffrey S. Skoll Chair of Technical Innovation and Entrepreneurship at the University of Toronto's Rotman School of Management and serves as chief economist in the Creative Destruction Lab.

PAUL IRVING is chairman of the Milken Institute Center for the Future of Aging, chairman of the board of Encore.org, and a distinguished scholar in residence at the University of Southern California, Davis School of Gerontology.

LESLIE K. JOHN is an associate professor at Harvard Business School.

MARISSA KIMSEY is a research associate at Harvard Business School.

KATRINA LAKE is the CEO of Stitch Fix.

ANDY NOBLE is a partner in Bain's Boston office, specializing in retail and organization.

NITIN NOHRIA is the dean of Harvard Business School.

ANNE-CLAIRE PACHE is the Chaired Professor in Philanthropy at ESSEC Business School.

MICHAEL E. PORTER is a University Professor at Harvard, based at Harvard Business School.

DARRELL K. RIGBY is a partner in the Boston office of Bain & Company. He heads the firm's global innovation and retail practices and is the author of *Winning in Turbulence*. (Harvard Business Review Press, 2009).

ERIN L. SCOTT is a senior lecturer in the Technological Innovation, Entrepreneurship, and Strategic Management Group at the MIT Sloan School of Management.

METIN SENGUL is an associate professor of strategy at Boston College's Carroll School of Management.

SCOTT STERN is the David Sarnoff Professor of Management at the MIT Sloan School of Management and serves as faculty director of the Martin Trust Center for MIT Entrepreneurship.

JEFF SUTHERLAND is a co-creator of the scrum form of agile innovation and the CEO of Scrum Inc., a consulting and training firm.

CATHERINE H. TINSLEY is the Raffini Family Professor of Management at Georgetown University's McDonough School of Business and the faculty director of the Georgetown University Women's Leadership Institute.

H. JAMES WILSON is a managing director of information technology and business research at Accenture Research. He

is a coauthor of *Human + Machine: Reimagining Work in the Age of AI* (Harvard Business Review Press, 2018)

Index

Invaluable insights
always at your fingertips

With an All-Access subscription to *Harvard Business Review*, you'll get so much more than a magazine.

Exclusive online content and tools you can put to use today

My Library, your personal workspace for sharing, saving, and organizing HBR.org articles and tools

---

Unlimited access to more than 4,000 articles in the *Harvard Business Review* archive

---

Subscribe today at hbr.org/subnow

The most important management ideas all in one place.

We hope you enjoyed this book from *Harvard Business Review*. Now you can get even more with HBR's 10 Must Reads Boxed Set. From books on leadership and strategy to managing yourself and others, this 6-book collection delivers articles on the most essential business topics to help you succeed.

HBR's 10 Must Reads Series

The definitive collection of ideas and best practices on our most sought-after topics from the best minds in business.

- Change Management
- Collaboration
- Communication
- Emotional Intelligence
- Innovation
- Leadership
- Making Smart Decisions
- Managing Across Cultures
- Managing People
- Managing Yourself
- Strategic Marketing
- Strategy
- Teams
- The Essentials

hbr.org/mustreads

Buy for your team, clients, or event.
Visit hbr.org/bulksales for quantity discount rates.

Harvard Business Review Press

HBR'S 10 MUST READS

BONUS ARTICLE

"The Feedback Fallacy"

By Marcus Buckingham and Ashley Goodall

The definitive management ideas of the year from **Harvard Business Review.**

2021

HBR's 10 Must Reads series is the definitive collection of ideas and best practices for aspiring and experienced leaders alike. These books offer essential reading selected from the pages of *Harvard Business Review* on topics critical to the success of every manager.

Titles include:

HBR's 10 Must Reads 2015
HBR's 10 Must Reads 2016
HBR's 10 Must Reads 2017
HBR's 10 Must Reads 2018
HBR's 10 Must Reads 2019
HBR's 10 Must Reads 2020
HBR's 10 Must Reads for CEOs
HBR's 10 Must Reads for New Managers
HBR's 10 Must Reads on AI, Analytics, and the New Machine Age
HBR's 10 Must Reads on Boards
HBR's 10 Must Reads on Building a Great Culture
HBR's 10 Must Reads on Business Model Innovation
HBR's 10 Must Reads on Change Management
HBR's 10 Must Reads on Collaboration
HBR's 10 Must Reads on Communication
HBR's 10 Must Reads on Design Thinking
HBR's 10 Must Reads on Diversity

HBR's 10 Must Reads on Emotional Intelligence
HBR's 10 Must Reads on Entrepreneurship and Startups
HBR's 10 Must Reads on Innovation
HBR's 10 Must Reads on Leadership
HBR's 10 Must Reads on Leadership (Vol. 2)
HBR's 10 Must Reads on Leadership for Healthcare
HBR's 10 Must Reads on Leadership Lessons from Sports
HBR's 10 Must Reads on Making Smart Decisions
HBR's 10 Must Reads on Managing Across Cultures
HBR's 10 Must Reads on Managing in a Downturn
HBR's 10 Must Reads on Managing People
HBR's 10 Must Reads on Managing People (Vol. 2)
HBR's 10 Must Reads on Managing Risk
HBR's 10 Must Reads on Managing Yourself
HBR's 10 Must Reads on Mental Toughness
HBR's 10 Must Reads on Negotiation
HBR's 10 Must Reads on Nonprofits and the Social Sectors
HBR's 10 Must Reads on Public Speaking and Presenting
HBR's 10 Must Reads on Reinventing HR
HBR's 10 Must Reads on Sales
HBR's 10 Must Reads on Strategic Marketing
HBR's 10 Must Reads on Strategy
HBR's 10 Must Reads on Strategy (Vol. 2)
HBR's 10 Must Reads on Strategy for Healthcare
HBR's 10 Must Reads on Teams
HBR's 10 Must Reads on Women and Leadership

HBR’s 10 Must Reads: The Essentials

HBR'S 10 MUST READS

The
Definitive Management Ideas of the Year from Harvard Business Review

2021

HARVARD BUSINESS REVIEW PRESS

Boston, Massachusetts

HBR Press Quantity Sales Discounts

Harvard Business Review Press titles are available at significant quantity discounts when purchased in bulk for client gifts, sales promotions, and premiums. Special editions, including books with corporate logos, customized covers, and letters from the company or CEO printed in the front matter, as well as excerpts of existing books, can also be created in large quantities for special needs.

For details and discount information for both print and ebook formats, contact booksales@harvardbusiness.org, tel. 800-988-0886, or www.hbr.org/bulksales.

Copyright 2020 Harvard Business School Publishing Corporation

All rights reserved

No part of this publication may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form, or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the prior permission of the publisher. Requests for permission should be directed to permissions@harvardbusiness.org, or mailed to Permissions, Harvard Business School Publishing, 60 Harvard Way, Boston, Massachusetts 02163.

First eBook Edition: June 2020

ISBN: 978-1-64782-003-9
eISBN: 978-1-64782-004-6

Contents

Editors' Note

Every year when our editorial group meets to assemble this collection, a few notable themes emerge. As we made our final selections, late last winter, the most prominent was the strange pace of change. Companies are adapting now to climate change—a threat that is both new and decades in the making, a danger to business and society that is already having direct consequences but whose ultimate effects have barely been felt. Similarly, the emergence of AI in business has been dizzying—but the changes it has wrought over the past half-decade may turn out to be minor compared with the tsunami of automation on the horizon. Inclusivity in the workplace has made great steps forward—companies are now considering how to provide a supportive environment for transgender people at work—even as the advancement of African-American leaders remains painfully slow. How are we to make sense of these swirling, contradictory changes? Is the practice of management improving faster than ever, or is it stuck in neutral?

A few weeks after we met, the coronavirus went worldwide. In what felt like an instant, all our notions about the pace of change and our predictions for the year 2021 were obliterated.

As of this writing, the world is many weeks into the first global pandemic lockdown. It's uncertain how much of this newly remote way of life is a passing moment and how much is the beginning of a new normal. Today parts of China, Europe, and the United States are beginning to ease restrictions, and we're looking ahead to the coming days with a mix of hope and fear.

Throughout, HBR's role has held steady. Our mission, "to improve the practice of management in a changing world," remains urgent. In recent weeks some forward-looking companies have shifted away from immediate crisis management and begun to reorganize for resilience and innovation in a time dominated by a desolate economic outlook, new habits, and employees who are hurting deeply. Effective management is needed more than ever to balance the existential imperatives of today with leadership for the long term.

Many pieces in this collection take on new meaning when seen through the lens of the pandemic. Cross-silo collaboration will be necessary not just to reinvent your own company but to lead the promising and innovative projects that will help the world rebound. Effective feedback and coaching become more challenging when workers are remote, lonely, and on the verge of burnout. Diversity and inclusion, along with fair hiring practices, are all the more pressing when so many jobs are at risk. The

challenges faced by dual-career couples and working parents become clearer every time a video meeting is interrupted by a Zoom-bombing toddler.

Meanwhile, the other starkly urgent issue of our time—climate change—looms larger than ever. If the pandemic has one positive outcome, it's the growth of our capacity to imagine what a collective disaster looks like. The veneer of invincibility on industry and humanity has dissolved. Covid-19 has put the consequences of inaction in plain sight and demonstrated that preventable yet irreparable damage can occur in *days*. We hope that business will be inspired to take brave action now on climate.

This book, our final selection of articles from the prepandemic era, points to a world ready for great change yet stuck in old habits and paradigms. The best HBR articles are enduring; as you read through them, look for practices that will help your business survive the crisis and reach new heights tomorrow.

For years managers have been encouraged to give candid feedback on just about everything workers do. But, as Marcus Buckingham and Ashley Goodall argue in **"The Feedback Fallacy,"** that doesn't actually help employees thrive, because identifying failure and telling people how to correct it will never produce great performance. Instead, when managers see a great outcome, they should acknowledge the person who created it and share their

impression of why it was a success. That fosters learning by showing what we're doing well. This was one of HBR's most popular—and widely debated—articles in many years.

Promising innovation and business opportunities require collaboration among functions, offices, and organizations. To realize them, companies must get people working together across boundaries, but that's a challenge for many leaders. **"Cross-Silo Leadership,"** by Tiziana Casciaro, Amy C. Edmondson, and Sujin Jang, explores the activities leaders can support to promote horizontal teamwork in their companies and help employees connect with and learn from people who think very differently from them.

At most large U.S. and multinational enterprises, diversity and inclusion have become imperatives, but the progress of African Americans remains slow. In **"Toward a Racially Just Workplace,"** Laura Morgan Roberts and Anthony J. Mayo argue that companies can take specific steps to achieve racial fairness. They can shift their focus from the most lucrative thing to the right thing, encourage open conversations about race, revamp diversity and inclusion programs to clarify goals and focus on proactive steps, and manage development across all career stages.

Thanks to technologies that enable constant, customized interactions, companies are building deeper ties with their customers. They can now address customer needs the

moment they arise—sometimes even earlier—and dramatically improve their customers' experiences, boost their own operational efficiencies, and gain competitive advantage. In **"The Age of Continuous Connection,"** Nicolaj Siggelkow and Christian Terwiesch identify four effective and connected strategies: filling customers' requests quickly and seamlessly; presenting individually tailored recommendations; reminding people of their needs and goals and nudging them to act; and anticipating what people want and delivering it without even being asked.

Innovative cultures are generally depicted as fun, characterized by a tolerance for failure and a willingness to experiment, and seen as psychologically safe, highly collaborative, and nonhierarchical. And research suggests that those qualities translate into better innovative performance. Even so, they're hard to create and sustain. Gary P. Pisano explains why in the McKinsey Award–winning **"The Hard Truth about Innovative Cultures."** Easy-to-like behaviors must be counterbalanced by some tougher ones: intolerance for incompetence, rigorous discipline, brutal candor, a high level of individual accountability, and strong leadership. Unless the tensions created by this paradox are carefully managed, attempts to create an innovative culture will fail.

Despite growing public awareness of the struggles that transgender individuals often face, many employers

remain ill-equipped to create supportive policies and workplace cultures. Christian N. Thoroughgood, Katina B. Sawyer, and Jennica R. Webster, the authors of **"Creating a Trans-Inclusive Workplace,"** demonstrate how companies can more effectively attract, retain, and promote the health and success of these workers through four areas of intervention: basic signs of trans inclusivity involving bathroom use, dress codes, and pronouns; effective support for gender transitions; trans-specific diversity trainings; and interventions to build resiliency.

Many executives assume that customer data can give you an unbeatable edge. The more customers you have, the more data you can gather, and analyzing that data allows you to offer a better product that attracts more customers. Although this thinking is usually wrong, Andrei Hagiu and Julian Wright show in **"When Data Creates Competitive Advantage"** that under the right conditions, customer data can help build competitive defenses. It all depends on whether the data offers high and lasting value, is proprietary, leads to improvements that can't be easily imitated, or generates insights that can be quickly incorporated. Only then will the virtuous cycle of data-enabled learning take off.

Companies now outsource much of the hiring process to vendors that in turn use subcontractors to scour LinkedIn and social media for potential candidates. When

applications come in, software sifts through them for key words that hiring managers want to see. According to Peter Cappelli in **"Your Approach to Hiring Is All Wrong,"** vendors offer an array of smart-sounding tools that claim to predict who will be a good hire—but it's yet to be proved that they produce satisfactory results. This article explores what's wrong with today's recruiting and hiring—failing to post jobs internally, advertising jobs that don't exist, misunderstanding the strengths and weaknesses of machine-learning models—and how to fix it.

In her research Jennifer Petriglieri, the author of **"How Dual-Career Couples Make It Work,"** found that such couples tend to experience three transitions when they are particularly vulnerable: when they first learn to work together as a couple; when they go through a midcareer or a midlife reinvention; and as they reach the final stages of their careers. Those who communicate during each transition about deeper work and personal issues such as values, boundaries, and fears have a better chance of emerging stronger from it, fulfilled both in their relationships and in their careers.

The use of artificial intelligence in business is no longer novel, yet most companies have run only ad hoc projects or applied AI in just a single business process. The key to capturing the full AI opportunity, according to Tim Fountaine, Brian McCarthy, and Tamim Saleh in **"Building**

the AI-Powered Organization," is to understand the organizational and cultural barriers faced by AI initiatives and work to lower them. That means shifting workers away from traditional mindsets—such as relying on top-down decision making—which often run counter to what's needed for AI. Leaders can also set up AI projects for success by conveying their urgency and benefits; investing heavily in AI education and adoption; and accounting for the company's AI maturity, business complexity, and innovation pace when deciding how work should be organized.

It's well past time to recognize that aggressive climate action is necessary if humanity is to survive and thrive. What's more, climate change is damaging the economy and company bottom lines *today*, not in some distant future. And although many large companies are cutting carbon emissions, their efforts are sadly inadequate given the scale of the crisis. In **"Leading a New Era of Climate Action,"** Andrew Winston recommends three approaches: using political influence to demand aggressive climate policies; empowering suppliers, customers, and employees to drive change; and rethinking investments and business models to eliminate waste and carbon. Companies must mobilize now to address this global threat.

Our final article is a special addition from hbr.org. On one of our first days of obligatory remote work during the

coronavirus pandemic, a few dozen of us found ourselves at home, looking at one another in a not-yet-commonplace grid of small, familiar faces. As in so many meetings in the early days of the crisis, we were talking about how we felt, and one colleague mentioned that she was experiencing grief. Heads nodded in all the panes. Editor Scott Berinato was inspired to reach out to David Kessler, a coauthor of *On Grief and Grieving: Finding the Meaning of Grief through the Five Stages of Loss*. Within a few days the interview **"That Discomfort You're Feeling Is Grief"** had reached an enormous worldwide audience and become the most-read article in the history of hbr.org. Kessler and Berinato discuss the stages of grief, techniques to make the feeling less intense, and the importance of a sixth stage—*meaning*—which can help us fathom painful events and become stronger as a result.

Finding the right words to open this book during the time of Covid-19 was a challenge. Books take a while to be published—so what is the right way to talk about an issue that's changing daily in a note that won't be read for months? Yet it gives us hope to know that you're still here with us and looking for ideas and inspiration from the best of HBR. We wish you, your family, and your teams good health during this surreal and scary time, and hope that our annual collection will help you and your business carve a bold path forward this year.

—The Editors

The Feedback Fallacy

by Marcus Buckingham and Ashley Goodall

THE DEBATE ABOUT FEEDBACK AT WORK ISN'T NEW. Since at least the middle of the last century, the question of how to get employees to improve has generated a good deal of opinion and research. But recently the discussion has taken on new intensity.

The ongoing experiment in "radical transparency" at Bridgewater Associates and the culture at Netflix, which the *Wall Street Journal* recently described as "encouraging harsh feedback" and subjecting workers to "intense and awkward" real-time 360s, are but two examples of the overriding belief that the way to increase performance in companies is through rigorous, frequent, candid, pervasive, and often critical feedback.

How should we give and receive feedback? we wonder. How much, and how often, and using which new app? And, given the hoopla over the approaches of Bridgewater and Netflix, how hard-edged and fearlessly candid should we be? Behind those questions, however, is another question that we're missing, and it's a crucial one. The search for ways to give and receive better feedback assumes that feedback is always useful. But the only

reason we're pursuing it is to help people do better. And when we examine *that*—asking, *How can we help each person thrive and excel?*—we find that the answers take us in a different direction.

To be clear, instruction—telling people what steps to follow or what factual knowledge they're lacking—can be truly useful: That's why we have checklists in airplane cockpits and, more recently, in operating rooms. There is indeed a right way for a nurse to give an injection safely, and if you as a novice nurse miss one of the steps, or if you're unaware of critical facts about a patient's condition, then someone should tell you. But the occasions when the actions or knowledge necessary to minimally perform a job can be objectively defined in advance are rare and becoming rarer. What we mean by "feedback" is very different. Feedback is about telling people what we think of their performance and how they should do it better—whether they're giving an effective presentation, leading a team, or creating a strategy. And on that, the research is clear: Telling people what we think of their performance doesn't help them thrive and excel, and telling people how we think they should improve actually *hinders* learning.

Underpinning the current conviction that feedback is an unalloyed good are three theories that we in the business world commonly accept as truths. The first is that other people are more aware than you are of your weaknesses, and that the best way to help you, therefore, is for them to show you what you cannot see for yourself. We can call this our *theory of the source of*

truth. You do not realize that your suit is shabby, that your presentation is boring, or that your voice is grating, so it is up to your colleagues to tell you as plainly as possible "where you stand." If they didn't, you would never know, and this would be bad.

The second belief is that the process of learning is like filling up an empty vessel: You lack certain abilities you need to acquire, so your colleagues should teach them to you. We can call this our *theory of learning.* If you're in sales, how can you possibly close deals if you don't learn the competency of "mirroring and matching" the prospect? If you're a teacher, how can you improve if you don't learn and practice the steps in the latest team-teaching technique or "flipped classroom" format? The thought is that you can't—and that you need feedback to develop the skills you're missing.

And the third belief is that great performance is universal, analyzable, and describable, and that once defined, it can be transferred from one person to another, regardless of who each individual is. Hence you can, with feedback about what excellence looks like, understand where you fall short of this ideal and then strive to remedy your shortcomings. We can call this our *theory of excellence.* If you're a manager, your boss might show you the company's supervisor-behaviors model, hold you up against it, and tell you what you need to do to more closely hew to it. If you aspire to lead, your firm might use a 360-degree feedback tool to measure you against its predefined leadership

competencies and then suggest various courses or experiences that will enable you to acquire the competencies that your results indicate you lack.

Idea in Brief

The Challenge

Managers today are bombarded with calls to give feedback—constantly, directly, and critically. But it turns out that telling people what we think of their performance and how they can do better is not the best way to help them excel and, in fact, can hinder development.

The Reality

Research shows that, first, we aren't the reliable raters of other people's performance that we think we are; second, criticism inhibits the brain's ability to learn; and, third, excellence is idiosyncratic, can't be defined in advance, and isn't the opposite of failure. Managers can't "correct" a person's way to excellence.

The Solution

Managers need to help their team members see what's working, stopping them with a "Yes! That!" and sharing their experience of what the person did well.

What these three theories have in common is self-centeredness: They take our own expertise and what we are sure is our colleagues' inexpertise as givens; they assume that my way is necessarily your way. But as it turns out, in extrapolating from what creates our own performance to what might create performance in others, we overreach.

Research reveals that none of these theories is true. The more we depend on them, and the more technology we base on them, the *less* learning and productivity we will get from others. To understand why and to see the path to a more effective way of improving performance, let's look more closely at each theory in turn.

The Source of Truth

The first problem with feedback is that humans are unreliable raters of other humans. Over the past 40 years psychometricians have shown in study after study that people don't have the objectivity to hold in their heads a stable definition of an abstract quality, such as *business acumen* or *assertiveness,* and then accurately evaluate someone else on it. Our evaluations are deeply colored by our own understanding of what we're rating others on, our own sense of what good looks like for a particular competency, our harshness or leniency as raters, and our own inherent and unconscious biases. This phenomenon is called the *idiosyncratic rater effect,* and it's large (more than half of your rating of someone else reflects your characteristics, not hers) and resilient (no training can lessen it). In other words, the research shows that feedback is more distortion than truth.

This is why, despite all the training available on how to *receive* feedback, it's such hard work: Recipients have to struggle through this forest of distortion in search of something that they recognize as themselves.

And because your feedback to others is always more you than them, it leads to systematic error, which is magnified when ratings are considered in aggregate. There are only two sorts of measurement error in the world: *random* error, which you can reduce by averaging many readings; and *systematic* error, which you can't. Unfortunately, we all seem to have left math class remembering the former and not the latter. We've built all our performance and leadership feedback tools as though assessment errors are random, and they're not. They're systematic.

Consider color blindness. If we ask a color-blind person to rate the redness of a particular rose, we won't trust his feedback—we know that he is incapable of seeing, let alone "rating," red. His error isn't random; it's predictable and explainable, and it stems from a flaw in his measurement system; hence, it's systematic. If we then decide to ask seven more color-blind people to rate the redness of our rose, their errors will be equally systematic, and averaging their ratings won't get us any closer to determining the actual redness of the rose. In fact, it's worse than this. Adding up all the inaccurate redness ratings—"gray," "pretty gray," "whitish gray," "muddy brown," and so on—and averaging them leads us *further away* both from learning anything reliable about the individuals' personal experiences of the rose and from the actual truth of how red our rose really is.

What the research has revealed is that we're all color-blind when it comes to abstract attributes, such as *strategic thinking*, *potential*, and *political savvy*. Our inability to rate others on them

is predictable and explainable—it is systematic. We cannot remove the error by adding more data inputs and averaging them out, and doing that actually makes the error bigger.

Worse still, although science has long since proven that we are color-blind, in the business world we assume we're clear-eyed. Deep down we don't think we make very many errors at all. We think we're reliable raters of others. We think we're a source of truth. We aren't. We're a source of error.

When a feedback instrument surveys eight colleagues about your business acumen, your score of 3.79 is far greater a distortion than if it simply surveyed one person about you—the 3.79 number is *all* noise, no signal. Given that (a) we're starting to see more of this sort of data-based feedback, (b) this data on you will likely be kept by your company for a very long time, and (c) it will be used to pay, promote, train, and deploy or fire you, you should be worried about just how fundamentally flawed it really is.

The only realm in which humans are an unimpeachable source of truth is that of their own feelings and experiences. Doctors have long known this. When they check up on you post-op, they'll ask, "On a scale of one to 10, with 10 being high, how would you rate your pain?" And if you say, "Five," the doctor may then prescribe all manner of treatments, but what she's unlikely to do is challenge you on your "five." It doesn't make sense, no matter how many operations she has done, to tell you your "five" is wrong, and that, actually, this morning your pain is

a “three.” It doesn’t make sense to try to parse what you mean by “five,” and whether any cultural differences might indicate that your “five” is not, in fact, a real “five.” It doesn’t make sense to hold calibration sessions with other doctors to ensure that your “five” is the same as the other “fives” in the rooms down the hall. Instead, she can be confident that you are the best judge of your pain and that all she can know for sure is that you will be feeling better when you rate your pain lower. Your rating is yours, not hers.

Just as your doctor doesn’t know the truth of your pain, we don’t know the truth about our colleagues, at least not in any objective way. You may read that workers today—especially Millennials—want to know where they stand. You may occasionally have team members ask you to tell them where they stand, objectively. You may feel that it’s your duty to try to answer these questions. But you can’t—none of us can. All we can do—and it’s not nothing—is share our own feelings and experiences, our own reactions. Thus we can tell someone whether his voice grates *on us*; whether he’s persuasive *to us*; whether his presentation is boring *to us.* We may not be able to tell him where he stands, but we can tell him where he stands *with us*. Those are our truths, not his. This is a humbler claim, but at least it’s accurate.

How We Learn

Another of our collective theories is that feedback contains useful information, and that this information is the magic ingredient that will accelerate someone's learning. Again, the research points in the opposite direction. Learning is less a function of adding something that isn't there than it is of recognizing, reinforcing, and refining what already is. There are two reasons for this.

The first is that, neurologically, we grow more in our areas of greater ability (our strengths are our development areas). The brain continues to develop throughout life, but each person's does so differently. Because of your genetic inheritance and the oddities of your early childhood environment, your brain's wiring is utterly unique. Some parts of it have tight thickets of synaptic connections, while others are far less dense, and these patterns are different from one person to the next. According to brain science, people grow far more neurons and synaptic connections where they already have the most neurons and synaptic connections. In other words, each brain grows most where it's already strongest. As Joseph LeDoux, a professor of neuroscience at New York University, memorably described it, "Added connections are therefore more like new buds on a branch rather than new branches." Through this lens, learning looks a lot like building, little by little, on the unique patterns already there within you. Which in turn means learning has to start by finding and understanding those patterns—your patterns, not someone else's.

Second, getting attention to our strengths from others catalyzes learning, whereas attention to our weaknesses smothers it. Neurological science also shows what happens to us when other people focus on what's working within us instead of remediating what isn't. In one experiment scientists split students into two groups. To one group they gave positive coaching, asking the students about their dreams and how they'd go about achieving them. The scientists probed the other group about homework and what the students thought they were doing wrong and needed to fix. While those conversations were happening, the scientists hooked each student up to a functional magnetic resonance imaging machine to see which parts of the brain were most activated in response to these different sorts of attention.

In the brains of the students asked about what they needed to correct, the sympathetic nervous system lit up. This is the "fight or flight" system, which mutes the other parts of the brain and allows us to focus only on the information most necessary to survive. Your brain responds to critical feedback as a threat and narrows its activity. The strong negative emotion produced by criticism "inhibits access to existing neural circuits and invokes cognitive, emotional, and perceptual impairment," psychology and business professor Richard Boyatzis said in summarizing the researchers' findings.

Focusing people on their shortcomings or gaps doesn't enable learning. It impairs it.

In the students who focused on their dreams and how they might achieve them, the sympathetic nervous system was not activated. What lit up instead was the parasympathetic nervous system, sometimes referred to as the "rest and digest" system. To quote Boyatzis again: "The parasympathetic nervous system … stimulates adult neurogenesis (i.e., growth of new neurons) …, a sense of well-being, better immune system functioning, and cognitive, emotional, and perceptual openness."

What findings such as these show us is, first, that learning happens when we see how we might do something better by adding some new nuance or expansion to our own understanding. Learning rests on our grasp of what we're doing well, not on what we're doing poorly, and certainly not on someone else's sense of what we're doing poorly. And second, that we learn most when someone else pays attention to what's working within us and asks us to cultivate it intelligently. We're often told that the key to learning is to get out of our comfort zones, but these findings contradict that particular chestnut: Take us very far out of our comfort zones, and our brains stop paying attention to anything other than surviving the experience. It's clear that we learn most in our comfort zones, because that's where our neural pathways are most concentrated. It's where we're most open to possibility, most creative, insightful, and productive. That's where feedback must meet us—in our moments of flow.

Excellence

We spend the bulk of our working lives pursuing excellence in the belief that while defining it is easy, the really hard part is codifying how we and everyone else on our team should get there. We've got it backward: Excellence in any endeavor is almost impossible to define, and yet getting there, for each of us, is relatively easy.

Excellence is idiosyncratic. Take funniness—the ability to make others laugh. If you watch early Steve Martin clips, you might land on the idea that excellence at it means strumming a banjo, waggling your knees, and wailing, "I'm a wild and crazy guy!" But watch Jerry Seinfeld, and you might conclude that it means talking about nothing in a slightly annoyed, exasperated tone. And if you watch Sarah Silverman, you might think to yourself, no, it's being caustic, blunt, and rude in an incongruously affectless way. At this point you may begin to perceive the truth that "funny" is inherent to the person.

Watch an NBA game, and you may think to yourself, "Yes, most of them are tall and athletic, but boy, not only does each player have a different role on the team, but even the players in the same role on the same team seem to do it differently." Examine something as specific and as limited as the free throws awarded after fouls, and you'll learn that not only do the top two free-throw shooters in history have utterly different styles, but one of them, Rick Barry—the best ever on the day he retired (look him up)—didn't even throw overhand.

Excellence seems to be inextricably and wonderfully intertwined with whoever demonstrates it. Each person's version of it is uniquely shaped and is an expression of that person's individuality. Which means that, for each of us, excellence is easy, in that it is a natural, fluid, and intelligent expression of our best extremes. It can be cultivated, but it's unforced.

Excellence is also not the opposite of failure. But in virtually all aspects of human endeavor, people assume that it is and that if they study what leads to pathological functioning and do the reverse—or replace what they found missing—they can create optimal functioning. That assumption is flawed. Study disease and you will learn a lot about disease and precious little about health. Eradicating depression will get you no closer to joy. Divorce is mute on the topic of happy marriage. Exit interviews with employees who leave tell you nothing about why others stay. If you study failure, you'll learn a lot about failure but nothing about how to achieve excellence. Excellence has its own pattern.

And it's even more problematic than that. Excellence and failure often have a lot in common. So if you study ineffective leaders and observe that they have big egos, and then argue that good leaders should not have big egos, you will lead people astray. Why? Because when you do personality assessments with highly effective leaders, you discover that they have very strong egos as well. Telling someone that you must lose your ego to be a good leader is flawed advice. Likewise, if you study poor

salespeople, discover that they take rejection personally, and then tell a budding salesperson to avoid doing the same, your advice will be misguided. Why? Because rigorous studies of the best salespeople reveal that they take rejection deeply personally, too.

As it happens, you find that effective leaders put their egos in the service of others, not themselves, and that effective salespeople take rejection personally because they are personally invested in the sale—but the point is that you will never find these things out by studying *ineffective* performance.

Since excellence is idiosyncratic and cannot be learned by studying failure, we can never help another person succeed by holding her performance up against a prefabricated model of excellence, giving her feedback on where she misses the model, and telling her to plug the gaps. That approach will only ever get her to adequate performance. Point out the grammatical flaws in an essay, ask the writer to fix the flaws, and while you may get an essay with good grammar, you won't get a piece of writing that transports the reader. Show a new teacher when her students lost interest and tell her what to do to fix this, and while you may now have a teacher whose students don't fall asleep in class, you won't have one whose students necessarily learn any more.

How to Help People Excel

If we continue to spend our time identifying failure as we see it and giving people feedback about how to avoid it, we'll languish

in the business of adequacy. To get into the excellence business we need some new techniques:

Look for outcomes

Excellence is an outcome, so take note of when a prospect leans into a sales pitch, a project runs smoothly, or an angry customer suddenly calms down. Then turn to the team member who created the outcome and say, “That! Yes, that!” By doing this, you’ll stop the flow of work for a moment and pull your colleague’s attention back toward something she just did that really worked.

There’s a story about how legendary Dallas Cowboys coach Tom Landry turned around his struggling team. While the other teams were reviewing missed tackles and dropped balls, Landry instead combed through footage of previous games and created for each player a highlight reel of when he had done something easily, naturally, and effectively. Landry reasoned that while the number of wrong ways to do something was infinite, the number of right ways, for any particular player, was not. It was knowable, and the best way to discover it was to look at plays where that person had done it excellently. From now on, he told each team member, “we only replay your winning plays.”

The Right Way to Help Colleagues Excel

IF YOU WANT to get into the excellence business, here are some examples of language to try.

Instead of	**Try**
Can I give you some feedback?	Here's my reaction.
Good job!	Here are three things that really worked for me. What was going through your mind when you did them?
Here's what you should do.	Here's what I would do.
Here's where you need to improve.	Here's what worked best for me, and here's why.
That didn't really work.	When you did *x*, I felt *y* or I didn't get that.
You need to improve your communication skills.	Here's exactly where you started to lose me.
You need to be more responsive.	When I don't hear from you, I worry that we're not on the same page.
You lack strategic thinking.	I'm struggling to understand your plan.
You should do *x* [in response to a request for advice].	What do you feel you're struggling with, and what have you done in the past that's worked in a similar situation?

Now on one level he was doing this to make his team members feel better about themselves because he knew the power of praise. But according to the story, Landry wasn't nearly as interested in praise as he was in learning. His instincts told him that each person would improve his performance most if he could see, in slow motion, what his own personal version of excellence looked like.

You can do the same. Whenever you see one of your people do something that worked for you, that rocked your world just a little, stop for a minute and highlight it. By helping your team member recognize what excellence looks like for her—by saying, “That! Yes, that!”—you’re offering her the chance to gain an insight; you’re highlighting a pattern that is already there within her so that she can recognize it, anchor it, re-create it, and refine it. That is learning.

Replay your instinctive reactions

Unlike Landry, you’re not going to be able to videotape your people. Instead, learn how to replay to them your own personal reactions. The key is not to tell someone how well she’s performed or how good she is. While simple praise isn’t a bad thing, you are by no means the authority on what objectively good performance is, and instinctively she knows this. Instead, describe what you experienced when her moment of excellence caught your attention. There’s nothing more believable and more authoritative than sharing what you saw from her and how it made you feel. Use phrases such as “This is how that came across for me,” or “This is what that made me think,” or even just “Did you see what you did there?” Those are your reactions—they are your truth—and when you relay them in specific detail, you aren’t judging or rating or fixing her; you’re simply reflecting to her the unique “dent” she just made in the world, as seen through your eyes. And precisely because it isn’t a judgment or a rating, it is at once more humble and more powerful.

On the flip side, if you're the team member, whenever your team leader catches you doing something right, ask her to pause and describe her reaction to you. If she says, "Good job!" ask, "Which bit? What did you see that seemed to work well?" Again, the point of this isn't to pile on the praise. The point is to explore the nature of excellence, and this is surely a better object for all the energy currently being pointed at "radical transparency" and the like. We're so close to our own performance that it's hard to get perspective on it and see its patterns and components. Ask for your leader's help in rendering the unconscious, conscious—so that you can understand it, improve at it, and, most important, do it again.

Never lose sight of your highest-priority interrupt

In computing, a high-priority interrupt happens when something requires a computer processor's immediate attention, and the machine halts normal operations and jumps the urgent issue to the head of the processing queue. Like computer processors, team leaders have quite a few things that demand their attention and force them to act. Many of them are problems. If you see something go off the rails—a poorly handled call, a missed meeting, a project gone awry—the instinct will kick in to stop everything to tell someone what she did wrong and what she needs to do to fix it. This instinct is by no means misguided: If your team member screws something up, you have to deal with it. But remember that when you do, you're merely remediating—and that remediating not only inhibits learning but also gets you

no closer to excellent performance. As we've seen, conjuring excellence from your team members requires a different focus from you. If you see somebody doing something that really works, stopping her and dissecting it with her isn't only a high-priority interrupt, it is your *highest*-priority interrupt. As you replay each small moment of excellence to your team member, you'll ease her into the "rest and digest" state of mind. Her understanding of what excellence looks and feels like within her will become more vivid, her brain will become more receptive to new information and will make connections to other inputs found in other regions of her brain, and she will learn and grow and get better.

Explore the present, past, and future

When people come to you asking for feedback on their performance or what they might need to fix to get promoted, try this:

Start with the *present*. If a team member approaches you with a problem, he's dealing with it *now*. He's feeling weak or challenged, and you have to address that. But rather than tackling the problem head-on, ask your colleague to tell you three things that are working for him *right now*. These things might be related to the situation or entirely separate. They might be significant or trivial. Just ask the question, and you're priming him with oxytocin—which is sometimes called the "love drug" but which here is better thought of as the "creativity drug." Getting him to think about specific things that are going well will

alter his brain chemistry so that he can be open to new solutions and new ways of thinking or acting.

Next, go to the *past.* Ask him, “When you had a problem like this in the past, what did you do that worked?” Much of our life happens in patterns, so it’s highly likely that he has encountered this problem at least a few times before. On one of those occasions he will almost certainly have found some way forward, some action or insight or connection that enabled him to move out of the mess. Get him thinking about that and seeing it in his mind’s eye: what he actually felt and did, and what happened next.

Finally, turn to the *future.* Ask your team member, “What do you already know you need to do? What do you already know works in this situation?” By all means offer up one or two of your own experiences to see if they might clarify his own. But operate under the assumption that he already knows the solution—you’re just helping him recognize it.

The emphasis here should not be on whys—“Why didn’t that work?” “Why do you think you should do that?”—because those lead both of you into a fuzzy world of conjecture and concepts. Instead, focus on the whats—“What do you actually want to have happen?” “What are a couple of actions you could take right now?” These sorts of questions yield concrete answers, in which your colleague can find his actual self doing actual things in the near-term future.

How to give people feedback is one of the hottest topics in business today. The arguments for radical candor and unvarnished and pervasive transparency have a swagger to them, almost as if to imply that only the finest and bravest of us can face these truths with nerveless self-assurance, that those of us who recoil at the thought of working in a climate of continual judgment are condemned to mediocrity, and that as leaders our ability to look our colleagues squarely in the eye and lay out their faults without blinking is a measure of our integrity.

But at best, this fetish with feedback is good only for correcting mistakes—in the rare cases where the right steps are known and can be evaluated objectively. And at worst, it's toxic, because what we want from our people—and from ourselves—is not, for the most part, tidy adherence to a procedure agreed upon in advance or, for that matter, the ability to expose one another's flaws. It's that people contribute their own unique and growing talents to a common good, when that good is ever-evolving, when we are, for all the right reasons, making it up as we go along. Feedback has nothing to offer to that.

We humans do not do well when someone whose intentions are unclear tells us where we stand, how good we "really" are, and what we must do to fix ourselves. We excel *only* when people who know us and care about us tell us what they experience and what they feel, and in particular when they see something within us that really works.

Originally published in March–April 2019. Reprint R1902G

Cross-Silo Leadership

by Tiziana Casciaro, Amy C. Edmondson, and Sujin Jang

THOUGH MOST EXECUTIVES recognize the importance of breaking down silos to help people collaborate across boundaries, they struggle to make it happen. That's understandable: It is devilishly difficult.

Think about your own relationships at work—the people you report to and those who report to you, for starters. Now consider the people in other functions, units, or geographies whose work touches yours in some way. Which relationships get prioritized in your day-to-day job?

We've posed that question to managers, engineers, salespeople, and consultants in companies around the world. The response we get is almost always the same: vertical relationships.

But when we ask, "Which relationships are most important for creating value for customers?" the answers flip. Today the vast majority of innovation and business-development opportunities lie in the interfaces between functions, offices, or organizations.

In short, the integrated solutions that most customers want—but companies wrestle with developing—require horizontal collaboration.

The value of horizontal teamwork is widely recognized. Employees who can reach outside their silos to find colleagues with complementary expertise learn more, sell more, and gain skills faster. Harvard's Heidi Gardner has found that firms with more cross-boundary collaboration achieve greater customer loyalty and higher margins. As innovation hinges more and more on interdisciplinary cooperation, digitalization transforms business at a breakneck pace, and globalization increasingly requires people to work across national borders, the demand for executives who can lead projects at interfaces keeps rising.

Our research and consulting work with hundreds of executives and managers in dozens of organizations confirms both the need for and the challenge of horizontal collaboration. "There's no doubt. We should focus on big projects that call for integration across practices," a partner in a global accounting firm told us. "That's where our greatest distinctive value is developed. But most of us confine ourselves to the smaller projects that we can handle within our practice areas. It's frustrating." A senior partner in a leading consulting firm put it slightly differently: "You know you should swim farther to catch a bigger fish, but it is a lot easier to swim in your own pond and catch a bunch of small fish."

One way to break down silos is to redesign the formal organizational structure. But that approach has limits: It's costly, confusing, and slow. Worse, every new structure solves some problems but creates others. That's why we've focused on identifying activities that facilitate boundary crossing. We've found that people can be trained to see and connect with pools of expertise throughout their organizations and to work better with colleagues who think very differently from them. The core challenges of operating effectively at interfaces are simple: *learning* about people on the other side and *relating* to them. But simple does not mean easy; human beings have always struggled to understand and relate to those who are different.

Leaders need to help people develop the capacity to overcome these challenges on both individual and organizational levels. That means providing training in and support for four practices that enable effective interface work.

1. Develop and Deploy Cultural Brokers

Fortunately, in most companies there are people who already excel at interface collaboration. They usually have experiences and relationships that span multiple sectors, functions, or domains and informally serve as links between them. We call these people *cultural brokers*. In studies involving more than 2,000 global teams, one of us—Sujin—found that diverse teams containing a cultural broker significantly outperformed diverse teams without one. (See "The Most Creative Teams Have a

Specific Type of Cultural Diversity," hbr.org, July 24, 2018.) Companies should identify these individuals and help them increase their impact.

Idea in Brief

The Challenge

Innovation initiatives, globalization, and digitalization increasingly require people to collaborate across functional and national boundaries. But breaking down silos remains frustratingly difficult.

The Cause

Employees don't know how to identify expertise outside their own work domains and struggle to understand the perspectives of colleagues who think very differently from them.

The Solution

Leaders can help employees connect with and relate to people across organizational divides by doing four things: developing and deploying "cultural brokers," who help groups overcome differences; encouraging and training workers to ask the right questions; getting people to see things through others' eyes; and broadening everyone's vision of networks of expertise inside and outside the company.

Cultural brokers promote cross-boundary work in one of two ways: by acting as a *bridge* or as an *adhesive*.

A bridge offers himself as a go-between, allowing people in different functions or geographies to collaborate with minimal disruption to their day-to-day routine. Bridges are most effective when they have considerable knowledge of both sides and can figure out what each one needs. This is why the champagne and

spirits distributor Moët Hennessy España hired two enologists, or wine experts, to help coordinate the work of its marketing and sales groups, which had a history of miscommunication and conflict. The enologists could relate to both groups equally: They could speak to marketers about the emotional content (the ephemeral "bouquet") of brands, while also providing pragmatic salespeople with details on the distinctive features of products they needed to win over retailers. Understanding both worlds, the enologists were able to communicate the rationale for each group's modus operandi to the other, allowing marketing and sales to work more synergistically even without directly interacting. This kind of cultural brokerage is efficient because it lets disparate parties work around differences without investing in learning the other side's perspective or changing how they work. It's especially valuable for one-off collaborations or when the company is under intense time pressure to deliver results.

Adhesives, in contrast, bring people together and help build mutual understanding and lasting relationships. Take one manager we spoke with at National Instruments, a global producer of automated test equipment. He frequently connects colleagues from different regions and functions. "I think of it as building up the relationships between them," he told us. "If a colleague needs to work with someone in another office or function, I would tell them, 'OK, here's the person to call.' Then I'd take the time to sit down and say, 'Well, let me tell you a little bit about how these guys work.'" Adhesives facilitate

collaboration by vouching for people and helping them decipher one another's language. Unlike bridges, adhesives develop others' capacity to work across a boundary in the future without their assistance.

Company leaders can build both bridging and adhesive capabilities in their organizations by hiring people with multifunctional or multicultural backgrounds who have the strong interpersonal skills needed to build rapport with multiple parties. Because it takes resilience to work with people across cultural divides, firms should also look for a *growth mindset*—the desire to learn and to take on challenges and "stretch" opportunities.

In addition, leaders can develop more brokers by giving people at all levels the chance to move into roles that expose them to multiple parts of the company. This, by the way, is good training for general managers and is what many rotational leadership-development programs aim to accomplish. Claudine Wolfe, the head of talent and development at the global insurer Chubb, maintains that the company's capacity to serve customers around the world rests on giving top performers opportunities to work in different geographies and cultivate an international mindset. "We give people their critical development experiences steeped in the job, in the region," she says. "They get coaching in the cultural norms and the language, but then they live it and internalize it. They go to the local bodega, take notice of the

products on the shelves, have conversations with the merchant, and learn what it really means to live in that environment."

Matrix organizational structures, in which people report to two (or more) groups, can also help develop cultural brokers. Despite their inherent challenges (they can be infuriatingly hard to navigate without strong leadership and accountability), matrices get people used to operating at interfaces.

We're not saying that everyone in your organization needs to be a full-fledged cultural broker. But consciously expanding the ranks of brokers and deploying them to grease the wheels of collaboration can go a long way.

2. Encourage People to Ask the Right Questions

It's nearly impossible to work across boundaries without asking a lot of questions. Inquiry is critical because what we see and take for granted on one side of an interface is not the same as what people experience on the other side.

Indeed, a study of more than 1,000 middle managers at a large bank that Tiziana conducted with Bill McEvily and Evelyn Zhang of the University of Toronto and Francesca Gino of Harvard Business School highlights the value of inquisitiveness in boundary-crossing work. It showed that managers with high levels of curiosity were more likely to build networks that spanned disconnected parts of the company.

But all of us are vulnerable to forgetting the crucial practice of asking questions as we move up the ladder. High-achieving

people in particular frequently fail to wonder what others are seeing. Worse, when we do recognize that we don't know something, we may avoid asking a question out of (misguided) fear that it will make us look incompetent or weak. "Not asking questions is a big mistake many professionals make," Norma Kraay, the managing partner of talent for Deloitte Canada, told us. "Expert advisers want to offer a solution. That's what they're trained to do."

Leaders can encourage inquiry in two important ways—and in the process help create an organization where it's psychologically safe to ask questions.

Be a role model

When leaders show interest in what others are seeing and thinking by asking questions, it has a stunning effect: It prompts people in their organizations to do the same.

Asking questions also conveys the humility that more and more business leaders and researchers are pointing to as vital to success. According to Laszlo Bock, Google's former senior vice president of people operations, humble people are better at bringing others together to solve tough problems. In a fast-changing business environment, humility—not to be confused with false modesty—is simply a strength. Its power comes from realism (as in *It really is a complex, challenging world out there; if we don't work together, we don't stand a chance*).

Gino says one way a leader can make employees feel comfortable asking questions is by openly acknowledging when

he or she doesn't know the answer. Another, she says, is by having days in which employees are explicitly encouraged to ask "Why?" "What if …?" and "How might we …?" (See "The Business Case for Curiosity," HBR, September–October 2018.)

Teach employees the art of inquiry

Training can help expand the range and frequency of questions employees ask and, according to Hal Gregersen of the MIT Leadership Center, can reinvigorate their sense of curiosity. But some questions are better than others. (See the exhibit "How to ask good questions.") And if you simply tell people to raise more questions, you might unleash interrogation tactics that inhibit rather than encourage the development of new perspectives. As MIT's Edgar Schein explains in his book *Humble Inquiry*, questions are the secret to productive work relationships—but they must be driven by genuine interest in understanding another's view.

How to ask good questions

Common pitfalls	Effective inquiry
Start with yes-or-no questions.	Start with open-ended questions that minimize preconceptions. ("How are things going on your end?" "What does your group see as the key opportunity in this space?")

Common pitfalls	Effective inquiry
Continue asking overly general questions ("What's on your mind?") that may invite long off-point responses.	As collaborations develop, ask questions that focus on specific issues but allow people plenty of room to elaborate. ("What do you know about *x*?" "Can you explain how that works?")
Assume that you've grasped what speakers intended.	Check your understanding by summarizing what you're hearing and asking explicitly for corrections or missing elements. ("Does that sound right—am I missing anything?" "Can you help me fill in the gaps?")
Assume the collaboration process will take care of itself.	Periodically take time to inquire into others' experiences of the process or relationship. ("How do you think the project is going?" "What could we do to work together more effectively?")

It's also important to learn how to request information in the least biased way possible. This means asking open-ended questions that minimize preconceptions, rather than yes-or-no questions. For instance, "What do you see as the key opportunity in this space?" will generate a richer dialogue than "Do you think this is the right opportunity to pursue?"

As collaborations move forward, it's helpful for team leaders or project managers to raise queries that encourage others to dive more deeply into specific issues and express related ideas or experiences. "What do you know about *x?"* and "Can you explain how that works?" are two examples. These questions are focused but neither limit responses nor invite long discourses that stray too far from the issue at hand.

How you process the answers also matters. It’s natural, as conversations unfold, to assume you understand what’s being said. But what people hear is biased by their expertise and experiences. So it’s important to train people to check whether they’re truly getting their colleagues’ meaning, by using language like “This is what I’m hearing—did I miss anything?” or “Can you help me fill in the gaps?” or “I think what you said means the project is on track. Is that correct?”

Finally, periodic temperature taking is needed to examine the collaborative process itself. The only way to find out how others are experiencing a project or relationship is by asking questions such as “How do you think the project is going?” and “What could we do to work together more effectively?”

3. Get People to See the World Through Others’ Eyes

Leaders shouldn’t just encourage employees to be curious about different groups and ask questions about their thinking and practices; they should also urge their people to actively consider others’ points of view. People from different organizational groups don’t see things the same way. Studies (including research on barriers to successful product innovation that the management professor Deborah Dougherty conducted at Wharton) consistently reveal that this leads to misunderstandings in interface work. It’s vital, therefore, to help people learn how to take the perspectives of others. One of us, Amy, has done research showing that ambitious cross-industry

innovation projects succeed when diverse participants discover how to do this. New Songdo, a project to build a city from scratch in South Korea that launched a decade ago, provides an instructive example. Early in the effort, project leaders brought together architects, engineers, planners, and environmental experts and helped them integrate their expertise in a carefully crafted learning process designed to break down barriers between disciplines. Today, in striking contrast to other "smart" city projects, New Songdo is 50% complete and has 30,000 residents, 33,000 jobs, and emissions that are 70% lower than those of other developments its size.

In a study of jazz bands and Broadway productions, Brian Uzzi of Northwestern University found that leaders of successful teams had an unusual ability to assume other people's viewpoints. These leaders could speak the multiple "languages" of their teammates. Other research has shown that when members of a diverse team proactively take the perspectives of others, it enhances the positive effect of information sharing and increases the team's creativity.

Creating a culture that fosters this kind of behavior is a senior leadership responsibility. Psychological research suggests that while most people are *capable* of taking others' perspectives, they are rarely *motivated* to do so. Leaders can provide some motivation by emphasizing to their teams how much the integration of diverse expertise enhances new value creation. But a couple of other tactics will help:

Organize cross-silo dialogues

Instead of holding one-way information sessions, leaders should set up cross-silo discussions that help employees see the world through the eyes of customers or colleagues in other parts of the company. The goal is to get everyone to share knowledge and work on synthesizing that diverse input into new solutions. This happens best in face-to-face meetings that are carefully structured to allow people time to listen to one another's thinking. Sometimes the process includes customers; one consulting firm we know started to replace traditional meetings, at which the firm conveyed information to clients, with a workshop format designed to explore questions and develop solutions in *collaboration* with them. The new format gives both the clients and the consultants a chance to learn from each other.

One of the more thoughtful uses of cross-silo dialogue is the "focused event analysis" (FEA) at Children's Minnesota. In an FEA people from the health system's different clinical and operational groups come together after a failure, such as the administration of the wrong medication to a patient. One at a time participants offer their take on what happened; the goal is to carefully document multiple perspectives *before* trying to identify a cause. Often participants are surprised to learn how people from other groups saw the incident. The assumption underlying the FEA is that most failures have not one root cause but many. Once the folks involved have a multifunctional picture

of the contributing factors, they can alter procedures and systems to prevent similar failures.

Hire for curiosity and empathy

You can boost your company's capacity to see the world from different perspectives by bringing on board people who relate to and sympathize with the feelings, thoughts, and attitudes of others. Southwest Airlines, which hires fewer than 2% of all applicants, selects people with empathy and enthusiasm for customer service, evaluating them through behavioral interviews ("Tell me about a time when …") and team interviews in which candidates are observed interacting.

4. Broaden Your Employees' Vision

You can't lead at the interfaces if you don't know where they are. Yet many organizations unwittingly encourage employees to never look beyond their own immediate environment, such as their function or business unit, and as a result miss out on potential insights employees could get if they scanned more-distant networks. Here are some ways that leaders can create opportunities for employees to widen their horizons, both within the company and beyond it:

Bring employees from diverse groups together on initiatives

As a rule, cross-functional teams give people across silos a chance to identify various kinds of expertise within their organization, map how they're connected or disconnected, and see how the

internal knowledge network can be linked to enable valuable collaboration.

At one global consulting firm, the leader of the digital health-care practice used to have its consultants speak just to clients' CIOs and CTOs. But she realized that that "unnecessarily limited the practice's ability to identify opportunities to serve clients beyond IT," she says. So she began to set up sessions with the entire C-suite at clients and brought in consultants from across all her firm's health-care practices—including systems redesign, operations excellence, strategy, and financing—to provide a more integrated look at the firm's health-care innovation expertise.

Those meetings allowed the consultants to discover the connections among the practices in the health-care division, identify the people best positioned to bridge the different practices, and see novel ways to combine the firm's various kinds of expertise to meet clients' needs. That helped the consultants spot value-generating opportunities for services at the interfaces between the practices. The new approach was so effective that, in short order, the leader was asked to head up a new practice that served as an interface across all the practices in the IT division so that she could replicate her success in other parts of the firm.

Urge employees to explore distant networks

Employees also need to be pushed to tap into expertise outside the company and even outside the industry. The domains of human knowledge span science, technology, business,

geography, politics, history, the arts, the humanities, and beyond, and any interface between them could hold new business opportunities. Consider the work of the innovation consultancy IDEO. By bringing design techniques from technology, science, and the arts to business, it has been able to create revolutionary products, like the first Apple mouse (which it developed from a Xerox PARC prototype into a commercial offering), and help companies in many industries embrace design thinking as an innovation strategy.

The tricky part is finding the domains most relevant to key business goals. Although many innovations have stemmed from what Abraham Flexner, the founding director of the Institute for Advanced Study, called "the usefulness of useless knowledge," businesses can ill afford to rely on open-ended exploratory search alone. To avoid this fate, leaders can take one of two approaches:

A *top-down approach* works when the knowledge domains with high potential for value creation have already been identified. For example, a partner in an accounting firm who sees machine learning as key to the profession's future might have an interested consultant or analyst in her practice take online courses or attend industry conferences about the technology and ask that person to come back with ideas about its implications. The partner might organize workshops in which the junior employee shares takeaways from the learning experiences and

brainstorms, with experienced colleagues, potential applications in the firm.

A *bottom-up approach* is better when leaders have trouble determining which outside domains the organization should connect with—a growing challenge given the speed at which new knowledge is being created. Increasingly, leaders must rely on employees to identify and forge connections with far-flung domains. One approach is to crowdsource ideas for promising interfaces—for example, by inviting employees to propose conferences in other industries they'd like to attend, courses on new skill sets they'd like to take, or domain experts they'd like to bring in for workshops. It's also critical to give employees the time and resources to scan external domains and build connections to them.

Breaking Down Silos

In today's economy everyone knows that finding new ways to combine an organization's diverse knowledge is a winning strategy for creating lasting value. But it doesn't happen unless employees have the opportunities and tools to work together productively across silos. To unleash the potential of horizontal collaboration, leaders must equip people to learn and to relate to one another across cultural and logistical divides. The four practices we've just described can help.

Not only is each one useful on its own in tackling the distinct challenges of interface work, but together these practices are

mutually enhancing: Engaging in one promotes competency in another. Deploying cultural brokers who build connections across groups gets people to ask questions and learn what employees in other groups are thinking. When people start asking better questions, they're immediately better positioned to understand others' perspectives and challenges. Seeing things from someone else's perspective—walking in his or her moccasins—in turn makes it easier to detect more pockets of knowledge. And network scanning illuminates interfaces where cultural brokers might be able to help groups collaborate effectively.

Over time these practices—none of which require advanced degrees or deep technical smarts—dissolve the barriers that make boundary-crossing work so difficult. When leaders create conditions that encourage and support these practices, collaboration across the interface will ultimately become second nature.

Originally published in May–June 2019. Reprint R1903J

Toward a Racially Just Workplace

by Laura Morgan Roberts and Anthony J. Mayo

“SUCCESS IS TO BE MEASURED not so much by the position that one has reached in life as by the obstacles which [one] has overcome while trying to succeed.”

Booker T. Washington, the educator, author, activist, and presidential adviser, wrote those words more than a century ago as a way of encouraging his African-American compatriots —many of them recently emancipated from slavery—to persist in the fight for equal rights and economic opportunities. He was proud of what he and his peers had achieved. He surely believed there was satisfaction in struggling against and surmounting bad odds. And yet we must also assume that he, along with millions of other freedom fighters, wanted future generations of black Americans to suffer fewer hardships. He hoped today’s black leaders would find easier paths to success.

Has that dream been realized? Having spent the past 20 years conducting and reviewing research on African-Americans’ advancement, particularly in the workplace, and having

collected our work and others' into a book, we must report that the answer is partly yes but mostly no.

No doubt, there has been progress. Civil rights laws have been passed and affirmed. Companies are committing to and investing heavily in diversity, because more corporate leaders acknowledge that it makes good business sense. And several black billionaires and CEOs sit on the respective ranking lists.

However, according to both quantitative and qualitative data, working African-Americans—from those laboring in factories and on shop floors to those setting C-suite strategy—still face obstacles to advancement that other minorities and white women don't. They are less likely than their white peers to be hired, developed, and promoted. And their lived experience at work is demonstrably worse even than that of other people of color.

These challenges might, as Washington said, make success sweeter for the few who overcome them. But a huge gap remains between what organizations are saying and doing to promote inclusion and the outcomes we're seeing for many black workers and managers. If leaders want to walk their talk, they must spearhead much more meaningful change. Instead of undervaluing and squandering black talent, they must recognize the resilience, robust sense of self, and growth mindset that, studies show, African-American people—as one of the most historically oppressed groups in the United States—bring to the table. They should work even harder to seek out

and support them, from entry-level recruitment to CEO succession.

We have not identified any major company that is doing this well on a broad scale. But research and lessons gleaned from other contexts can point the way forward. In our work with leading management thinkers and practitioners across the country, we have arrived at a four-step strategy to help companies move toward greater and better representation for black leaders. It involves shifting from an exclusive focus on the business case for racial diversity to embracing the moral one, promoting real conversations about race, revamping diversity and inclusion programs, and better managing career development at every stage. Given the increasing importance of purpose and social impact to employees, customers, and other stakeholders, we believe there's no better time to make this transformation. We also believe our framework can be adapted for other marginalized groups in the United States and around the world.

Taking these steps won't be easy; executives will need to think deeply about their ethics and corporate culture and exert extra effort for a cause they may not consider central to their business. But the reward will be great: maximizing the human potential of everyone in the workplace.

Idea in Brief

At most large U.S. and multinational organizations, diversity and inclusion have become imperatives, but African-Americans' progress remains slow. Even worse is the lived experience of black employees, who often feel like outsiders—and are tempted to walk out the door. Laura Morgan Roberts and Anthony J. Mayo argue that organizations should take specific steps toward racial justice:

- Shift from an exclusive focus on the most lucrative thing to do to the right thing to do.
- Encourage open conversations about race.
- Revamp diversity and inclusion programs to clarify goals and focus on proactive steps.
- Manage career development across all life stages, from campus recruitment to the consideration of black executives for top jobs.

These steps won't be easy, but maximizing the human potential of everyone in the workplace is the ultimate reward.

Underrepresented, Unsupported, Unfulfilled

At most large U.S. and multinational organizations, diversity and inclusion (D&I) has become an imperative. Companies are pushing for minority recruitment, paying for antibias training, and sponsoring nonwhite employees for high-potential leadership-development programs. Research has shown, and a great many executives now understand, that a heterogeneous workforce yields more innovation and better performance than a homogeneous one does.

And yet 55 years after the passage of the Civil Rights Act and decades into these corporate D&I efforts, African-Americans' progress toward top management roles and greater economic well-being and influence remains slow to nonexistent. Let's look first at the demographics.

The Big Idea: Advancing Black Leaders

"Toward a Racially Just Workplace" is the lead article of HBR's **The Big Idea: Advancing Black Leaders.** You can read the rest of the series at hbr.org/blackleaders:

- "The Costs of Code-Switching," by Courtney L. McCluney, Kathrina Robotham, Serenity Lee, Richard Smith, and Myles Durkee
- "The Day-to-Day Work of Diversity and Inclusion," by Paige Cohen and Gretchen Gavett
- "Why So Many Organizations Stay White," by Victor Ray
- "Success Comes from Affirming Your Potential," by Laura Morgan Roberts and Anthony J. Mayo

What the numbers say

Yes, we can point to the rise of several prominent black leaders, from media figures Oprah Winfrey, Robert Johnson, and Jay-Z to financiers Ken Chenault and Robert Smith and sports-stars-turned-businesspeople Serena Williams, Michael Jordan, and

LeBron James. Most notably, America elected its first African-descended president, Barack Obama, in 2008 and reelected him in 2012. The number of African-Americans earning bachelor's and graduate degrees continues to increase. And black people account for 12% of the U.S. workforce, close to their 13.4% representation in the general population.

However, in the words of leaders from the Toigo Foundation, a career advancement organization for underrepresented groups, such evidence merely gives us "the illusion of inclusion." In fact, research shows that in the United States, the wealth gap between blacks and others continues to widen; experts predict that black families' median wealth will decrease to $0 by 2050, while that of white families will exceed $100,000. Just 8% of managers and 3.8% of CEOs are black. In the *Fortune* 500 companies, there are currently only three black chief executives, down from a high of 12 in 2002. And at the 16 *Fortune* 500 companies that report detailed demographic data on senior executives and board members, white men account for 85% of those roles.

Black leaders have struggled to make inroads in a variety of influential industries and sectors. At U.S. finance companies, only 2.4% of executive committee members, 1.4% of managing directors, and 1.4% of senior portfolio managers are black. A mere 1.9% of tech executives and 5.3% of tech professionals are African-American. Black representatives and senators account for 9% of the U.S. Congress. The average black partnership rate

at U.S. law firms from 2005 to 2016 was 1.8%. Only 7% of U.S. higher education administrators and 8% of nonprofit leaders are black. And just 10% of U.S. businesses are owned by black men and women. As the Toigo Foundation points out, all this has a cascading impact on economic development, housing, jobs, quality of schools and other services, access to education, infrastructure spending, consumer credit, retirement savings, and more.

What it's like at work

Underrepresentation is bad enough. But even worse, according to extensive research, is the lived experience of black employees and managers in the U.S. workplace. African-Americans continue to face both explicit racism—stoked by the rise of white nationalism in the past few years—and subtle racism on the job. In the latter category, University of Utah professor emeritus Arthur Brief points to "aversive" racism (when people avoid those of different races or change their behavior around them) along with "modern" racism (when people believe that because blacks can now compete in the marketplace, they no longer face discrimination). Microagressions—for example, when a white male visitor to an office assumes that a black female executive is a secretary—are also common.

Although companies claim they want to overcome these explicit and implicit biases and hire and promote diverse

candidates, they rarely do so in effective ways. When Harvard Business School's emeriti professors David A. Thomas and John Gabarro conducted an in-depth six-year study of leaders in three companies, they found that people of color had to manage their careers more strategically than their white peers did and to prove greater competence before winning promotions. And research by Lynn Perry Wooten, the dean of Cornell University's Dyson School, and Erika Hayes James, the dean of Emory University's Goizueta Business School, shows that black leaders who do rise to the top are disproportionately handed "glass cliff" assignments, which offer nice rewards but carry a greater risk of failure. Other research, such as Duke University professor Ashleigh Rosette's studies of black leaders, has shown widespread racial differences in hiring, performance ratings, promotions, and other outcomes.

There is also an emotional tax associated with being black in the American workplace. Research by the University of Virginia's Courtney McCluney and Catalyst's Dnika Travis and Jennifer Thorpe-Moscon shows that because black employees feel a heightened sense of difference among their mostly white peers, their ability to contribute is diminished. "The sense of isolation, of solitude, can take a toll," one leader told them. "It's like facing each day with a core of uncertainty … wondering … if the floor you're standing on is concrete or dirt … solid or not."

Many black professionals have reported to Toigo that they are expected to be “cultural ambassadors” who address the needs of other black employees, which leaves them doing two jobs: “the official one the person was hired to do, and a second one as champion for members of the person’s minority group,” as one put it. Across industries, sectors, and functions, they also experience the “diversity fatigue” that arises from constantly engaging in task forces, trainings, and conversations about race as they are tapped to represent their demographic.

And black leaders in particular struggle with feeling inauthentic at work. Research by McGill University’s Patricia Faison Hewlin shows that many minorities feel pressured to create “facades of conformity,” suppressing their personal values, views, and attributes to fit in with organizational ones. But as Hewlin and her colleague Anna-Maria Broomes found in a survey of 2,226 workers in various industries and corporate settings, African-Americans create these facades more frequently than other minority groups do and feel the inauthenticity more deeply. They might chemically relax (straighten) their hair, conform with coworkers’ behavior, “whitewash” their résumés by deleting ethnic-sounding names or companies, hide minority beliefs, and suppress emotions related to workplace racism.

As a result of all the above, black workers feel less supported, engaged, and committed to their jobs than their nonblack peers do, as research from Georgetown University’s Ella

Washington, Gallup's Ellyn Maese and Shane McFeely, and others has documented. Black managers report receiving less psychosocial support than their white counterparts do. Black employees are less likely than whites or Hispanics to say that their company's mission or purpose makes them feel that their job is important, that their coworkers will do quality work, and that they have opportunities to learn and grow. Black leaders are more likely than white ones to leave their organizations. It's clear that the norms and cultural defaults of leadership in most organizations create an inhospitable environment that leaves even those black employees who have advanced feeling like outsiders—and in some cases pushes them out the door.

Relatively high pay and impressive pedigrees don't help much: According to a survey of diverse professionals with bachelor's or graduate degrees and average annual incomes of $100,000 or more that one of us (Laura) conducted with colleagues at the Partnership, a nonprofit organization specializing in diversity and leadership development, African-Americans report the lowest levels of both manager and coworker support, commitment, and job fit and the highest levels of feeling inauthentic and wanting to leave their jobs. Studies of black Harvard Business School and Harvard Law School graduates have similarly found that matriculating from highly respected institutions does not shield one from obstacles. When surveyed years and even decades after graduating, black Harvard MBAs expressed less satisfaction

than their white counterparts with opportunities to do meaningful work, to realize professional accomplishments, and to combine career with personal and family life. “Perhaps it sounds naive, but [coming out of HBS] I did not expect race to have any bearing in my career,” one told us. “I was wrong.”

Leading Change

As we said earlier, diversity and inclusion efforts have been gaining traction, and workforces are becoming increasingly multiracial. But given the dearth of black leaders, we would like to see companies jump-start their efforts in four ways.

First, move away from the business case and toward a moral one
The dozens of D&I executives we talked to in the course of our research tell us they sometimes feel they’ve taken the business case for diversity as far as it can go. When Weber Shandwick surveyed 500 chief diversity officers at companies with revenue of $500 million or more, results confirmed that proving that ROI—showing that inclusive teams yield more-creative ideas that appeal to broader customer bases, open new markets, and ultimately drive better performance—is one of the biggest challenges.

The research on this is clear. A 2015 McKinsey report on 366 public companies found that those in the top quartile for ethnic and racial diversity in management were 35% more likely than others to have financial returns above the industry mean.

Various studies have shown that teams composed of both white and black people are more likely to focus on facts, carefully process information, and spur innovation when the organizational culture and leadership support learning across differences.

With the right knowledge, skills, and experience, African-American employees and managers can add as much business value as anyone else. They may have greater insights about creating and selling offerings for minority consumer groups that end up appealing to white consumers as well. As one of us (Tony) showed in research with Nitin Nohria, now the dean of Harvard Business School, and Eckerd College's Laura Singleton, some of the most successful black entrepreneurs are those who —in some cases *because* they were marginalized—built companies to serve their same-race peers, particularly in the personal care, media, and fashion arenas. Examples include the 19th-century black-hair-care trailblazer Madam C. J. Walker, Black Entertainment Television's Robert and Sheila Johnson, and Daymond John, who launched the FUBU clothing line.

So, experts agree that diversity enhances business outcomes when managed well. But given the limited progress African-Americans have made in most of corporate America, it seems clear that the sound business arguments for inclusion are not enough. At many companies, D&I executives still struggle for airtime in the C-suite and for resources that can move their organizations beyond the tokenism of, say, one black executive

in the senior ranks. Their business cases don't appear to have been as persuasive as those presented by their marketing, operations, and accounting colleagues, which have a more direct effect on the bottom line.

And in more-progressive companies—ones truly committed to inclusion—a different kind of pushback sometimes occurs: If a team incorporates women, Asians, Latinos, and representatives of the LGBTQ community alongside white men, if it has data geeks and creative types, extroverts and introverts, Harvard MBAs and college dropouts, able-bodied and physically challenged members, isn't it diverse enough? Our answer: not when teams, especially those at the highest levels, leave out the most marginalized group in the United States.

Thus we turn to the moral case. Many in the U.S. business community have begun to push for a more purpose-driven capitalism that focuses not just on shareholder value but also on shared value—benefits that extend to employees, customers, suppliers, and communities. This movement, toward what the University of Toronto's Sarah Kaplan calls the 360° Corporation, wants corporate leaders to consider both the financial and the ethical implications of all their decisions. We believe that one of its pillars should be proportionate representation and wages for black Americans.

Why this group in particular? As the *New York Times*'s excellent 1619 Project highlighted, we are exactly four

centuries away from the start of slavery—the kidnapping, forced labor, mistreatment, and often murder of African people —in the United States. And we are just 154 years away from its end. Although discrimination based on race and other factors was outlawed by the Civil Rights Act of 1964, the effects of slavery and the decades of discrimination and disenfranchisement that followed it continue to hold back many descendants of enslaved people (and those from different circumstances who have the same skin color). Alarmingly, racism and racist incidents are on the rise: According to the FBI, the number of hate crimes committed in the United States rose by 17% from 2016 to 2017, marking the third consecutive year of increases.

We also can't forget that a compelling business case can be—and has been—made for all the atrocities listed above. Indeed, when invoked absent humanistic and ethical principles, a "business case" has legitimated exploitative actions throughout history. White landowners argued that the economic welfare of the colonies and the health of a young country depended on keeping black people in chains. White business owners in the Jim Crow South and segregated neighborhoods across the country claimed that sales would suffer if black customers and residents—who in the absence of land and good jobs had amassed little wealth—were allowed in, because that would turn rich white customers away. And white executives have long benefited because people of color with

less access to high-quality education and high-wage employment were forced into low-paying commercial and household jobs, from coal mining and call center work to cleaning, cooking, and caregiving.

So the case for racial diversity and the advancement of African-Americans can't be solely about increasing innovation or providing access to and legitimacy in minority markets to maximize revenue and profits. We can't simply ask, "What's the most lucrative thing to do?" We must also ask, "What's the right thing to do?" The imperative should be creating a context in which people of all colors, but especially those who have historically been oppressed, can realize their full potential. This will involve exploring and understanding the racist history that has shaped various groups' access to resources and opportunities and that undergirds contemporary bias. It means emphasizing equity and justice.

How might this work? Starbucks has made some attempts. In the wake of protests following the 2014 fatal shooting of Michael Brown by police in Ferguson, Missouri, the coffee chain announced RaceTogether, which aimed to spark a national conversation about race relations by having baristas write that phrase on customers' cups. The campaign fell flat because it was perceived more as a profit-minded marketing stunt than as a good-faith effort to change the status quo. Subsequent initiatives, perhaps designed with ethics more squarely in mind, have garnered a more positive response. In 2015

Starbucks launched a hiring program to recruit disadvantaged youths, including African-Americans; in 2017 it expanded that program and added one to recruit refugees; and after a racially charged incident at one of its cafés in 2018, it closed all its U.S. cafés for a day of employee antibias training. Consider, too, Nike's decision to launch a marketing campaign headlined by Colin Kaepernick, the NFL quarterback who failed to get picked up by a team after he began kneeling during the national anthem to protest the unfair treatment of African-Americans. The campaign created a backlash among anti-Kaepernick consumers and a #BoycottNike hashtag, but the sports apparel brand stood by its tagline: "Believe in something. Even if it means sacrificing everything." We applaud these steps and hope organizations will go even further in learning how to practice racial inclusion in their workplaces.

Some organizations have invoked the moral case for action in other contexts. Think of how Patagonia supports environmental protections by committing to donate either 1% of sales or 10% of profits (whichever is larger) to advocacy groups. And recall that Dick's Sporting Goods pulled assault weapons and high-capacity magazines from its stores following the Parkland, Florida, school shooting, even though it projected—accurately—that the move would mean a $250 million hit to sales. (It's important to note that over the long term, none of those companies suffered from their choices.)

Such stances take courage. But by combining the business case and the moral one, leaders can make a more powerful argument for supporting black advancement.

Second, encourage open conversations about race

As Dartmouth College's Ella Bell and the University of Pretoria's Stella Nkomo note in the introduction to our book, "Organizations are in society, not apart from it." And although President Obama's election brought some talk of a post-racial era in the United States, the stories and statistics that have come out in the past few years show that racism still exists, which means that race still matters and needs to be discussed, candidly and frequently, in the workplace.

Those conversations will not immediately feel comfortable. Research shows that although many people are happy to talk about "diversity" or "inclusion," their enthusiasm drops significantly when the subject is "race." Most of us don't like to think very hard about where minorities sit and what power they wield (or don't) within our organizations—much less discuss it. When we examine who has been excluded in what ways over what period of time, the concept of white privilege might come up. And majority-group employees might express concerns about reverse discrimination. (According to an Ernst & Young study of 1,000 U.S. workers, one-third of respondents said that a corporate focus on diversity has overlooked white men.) Charged topics like these can provoke resentment, anger,

and shame. But we need real exchanges about them if we want to dispel the notion that corporations are pure meritocracies and to ensure that everyone feels heard, supported, and authentic at work.

Senior leaders—most of whom are white men—must set the tone. Why? In one survey, nearly 40% of black employees said they feel it is *never* acceptable to speak out about experiences of bias—a silence that can become corrosive. Another study showed that among black professionals who aspire to senior leadership positions, the most frequently adopted strategy is to avoid talking about race or other issues of inequality, for fear of being labeled an agitator. Other research has indicated that the only CEOs and lower-level managers not penalized for championing diversity are white men.

To create a culture of psychological safety and pave the way for open communication will require a top-down directive and modeling through informal and formal discussions in which people are asked to share ideas, ask questions, and address issues without fear of reprisal. Managers down the line will need training in encouraging and guiding such exchanges, including inviting black employees and leaders to share their experiences—the good, the bad, and the ugly. Participants should be trained to prepare for such conversations by reflecting on their own identities and the comments and situations that trigger strong emotions in them. As detailed by Columbia University's Valerie Purdie-Greenaway and the

University of Virginia's Martin Davidson, the goal is to shift the entire organization to a racial-learning orientation.

Again, a movement from another context—#MeToo—sheds light on how to do so. Revelations of abuse and harassment and the outpouring of women's stories that followed, many about incidents that happened in the workplace, forced corporate leaders to focus on those issues. Bad actors were fired, women felt empowered to speak up, and awareness of gender discrimination increased. Although #BlackLivesMatter has had similar success highlighting and sparking discussions around police brutality, there is no #BlackLivesAtWork. There should be.

We see some positive signs on this front. Over the past few years several prominent leaders, including PwC's Tim Ryan, Interpublic Group's Michael Roth, Kaiser Permanente's Bernard Tyson, and AT&T's Randall Stephenson, have initiated companywide discussions of race. For example, PwC brought in Mellody Hobson, president and co-CEO of Ariel Investments and a prominent African-American leader, to talk to employees about being "color-brave" instead of "color-blind" at work, and it has offered guides for continuing the discussion. At Morgan Stanley, global head of D&I Susan Reid has promoted intimate conversations about race in networking groups and an hour-long forum on race in the current social climate. The latter was moderated by the company's vice chairman and featured its chief marketing officer, its head of prime brokerage, and a

Fortune reporter who covers racial issues; it was attended by 1,500 employees, and videos of the event were shared across the firm. Greenaway and Davidson also point to a mostly white male financial services firm that instituted Know Us, a program of small-group cross-race dialogues on racially relevant topics.

Over time these conversations will start to happen informally and organically in groups and among individuals at all levels of an organization, deepening interpersonal cross-race relationships. In one consulting company cited by Greenaway and Davidson, nonblack employees started a book club open to all but focused on black writers; the group has visited African-American museums and historical sites. One-on-one interactions can be even more meaningful, as the psychologist colleagues Karen Samuels (who is white) and Kathryn Fraser (who is black) describe. “It was important to name our racial and cultural differences and to examine how my perspective was naive regarding her reality,” Samuels explains.

Third, revamp D&I programs

Any corporate diversity and inclusion program is better than none, but most that exist today are not designed to sustain a focus on racial equity. Many are siloed within the HR department, lack C-suite support, or are given to women or people of color to manage in addition to their day jobs. Some are more show than go, resting on philosophical statements about inclusion rather than outlining concrete steps for

advancing nonwhites. Others limit their efforts to antibias and cultural competence training—preempting problems but, again, not propelling anyone forward. Most take a broad-brush approach to diversity, attempting to serve all minorities plus white women, LGBTQ employees, and those who are neurodiverse or disabled and offering uniform training and leadership development that ignore historical patterns of exclusion, marginality, and disadvantage for each group. They might focus too heavily on recruitment and retention—filling the pipeline and high-potential groups with black employees but failing to support them past middle-management roles. Most troubling, as Courtney McCluney and San Francisco State University's Verónica Rabelo have shown, a significant portion of D&I programs try to "manage blackness"—that is, impose "desirable" and "professional" (read: white) norms and expectations on rising African-American stars, thus preserving rather than shifting the status quo. They train black executives to fit into the existing organizational culture rather than encourage them to broaden it by bringing their true and most productive selves to work.

How can we improve such programs? By tackling their shortcomings one by one. Here are several steps organizations can take.

- Give D&I sustained C-suite support and recognize and reward the people who contribute to its initiatives—for

example, by having your chief diversity officer report directly to the CEO and tracking inclusion initiative participation in performance reviews and promotion and pay raise discussions.

- Equip and invite white men to take up the mantle—say, by bringing them into D&I programs and assigning some of them to leadership roles.
- Challenge those running D&I efforts to set clear goals for how representation, organizational networks, and access to resources should change across functions and levels over time and how black employees' perceptions, engagement, and well-being should improve, and then measure the efforts' effectiveness with data analysis and qualitative surveys.
- Shift from preventative measures, such as antibias training, to proactive ones, such as upping the number of black candidates considered for open positions and stretch roles.
- Abandon one-size-fits-all and color-blind leadership-development practices in favor of courses and coaching tailored to specific groups—or better yet, adopt personalized plans that recognize the multifaceted nature of each individual.
- Help black employees and rising leaders throughout their careers, including teaching managers the skills they need

to support D&I efforts.

- Stop asking black employees to blend in; instead, emphasize the value of a workplace that embraces all styles and behaviors.

In sum, D&I needs to be an ethos that permeates the entire organization, championed not just by the HR department but by everyone, and especially managers, so that its importance is clear. The Toigo Foundation's leaders draw a parallel between this idea and the total quality management movement of the 1980s, which, with top-down support and the establishment of key performance indicators, became a pervasive way of working and thinking that filtered down to every function and level.

Few companies to date have taken diversity and inclusion that far. But some are moving in the right direction, including JPMorgan Chase, which in 2016 launched a board- and CEO-supported Advancing Black Leaders strategy—staffed and managed separately from other D&I initiatives—focused on filling the firm's pipeline with black talent and retaining and promoting those workers. SAP's Black Employee Network helped launch its partnership with Delaware State University through Project Propel, which offers tech training and skills development to students from historically black colleges and universities (HBCUs), with the goal of building an employee pipeline. The Network also encouraged SAP to sponsor Silicon

Valley's Culture Shifting Weekend, which brings together more than 200 African-American and Hispanic executives, entrepreneurs, innovators, and social impact leaders to discuss diversifying the tech industry. Pfizer tracks numerous D&I metrics and notes that 21% of its workforce—21,000 people—are actively involved in its D&I efforts.

Finally, manage career development across all life stages

African-Americans today are securing good university educations in record numbers. HBCUs, in particular, create a sizable pipeline of young talent for organizations to tap into. Companies can, of course, step up their campus recruiting efforts, but efforts to advance black leaders must extend far beyond that.

If more African-Americans are to rise through the ranks, robust—and careful—investment in retention and development is required. Research by the University of Georgia's Kecia Thomas and colleagues has shown that many black women get this kind of support early in their careers, but it comes with a price: They are treated like "pets" whom white leaders are happy to groom, but the further they progress, the more that favored status begins to undermine them. Those who reject the pet identity, meanwhile, are perceived as threatening and face hostility and distancing from coworkers.

Mentoring is useful, and our study of black HBS graduates shows that they were more likely than their white peers to

have been formally assigned to mentors. But they derived less value from the relationship and said that informal mentorship—having senior executives (white or minority) connect with them naturally through work groups or common interests—was more effective. “A mentor helps you navigate the power structure of the firm, especially when there is no one in senior management who looks like you,” one study participant told us.

Early in their careers, black employees need safe spaces to grow and develop and to experience authentic failures and successes without being subsumed in narratives of racial limitation. Managers and mentors can provide the necessary cover. We found that the black Harvard MBAs who did reach top management positions (13% of women, 19% of men) had been bolstered by networks of supporters.

Sponsorship—that is, recommending black employees for promotions and stretch assignments—is even more important. Other key factors that have propelled black Harvard MBAs into senior executive roles are line or general management experience and global assignments. With many qualified and ambitious people vying for such opportunities, politics often plays a role. So African-Americans need more influential people in their corners, pressing their cases to decision-makers.

Candid feedback early on is also critical. This doesn’t mean pushing protégés to assimilate (to look and act “more white”); as we’ve shown, that’s counterproductive. It should focus on

identifying and enhancing their unique strengths, overcoming skill or knowledge weaknesses, and positioning them to realize their full potential.

At later stages of their careers, black executives should be seriously considered for high-stakes and high-profile positions and supported in the pursuit of outside interests, such as board seats, that enhance visibility. And while taking care not to tokenize but rather to create opportunities for multiple candidates, organizations can highlight those executives as role models who redefine norms of leadership and can encourage them to pass that baton by transferring connections and endorsements, sharing wisdom through storytelling, and creating opportunities for the next generation to assume senior roles. Needs differ by career stage, a fact that most published models of diversity and inclusion do not address but that is embedded in impactful programs such as the Toigo Foundation, the Partnership, and the Executive Leadership Council.

Despite antidiscrimination laws and increasing corporate investment in diversity efforts, race continues to be a major barrier to advancement in the U.S. workplace. We are far from realizing the principles of equal opportunity and meritocracy. Rather than looking to the few black leaders who have succeeded as exemplars of exceptionalism who have beaten

almost insurmountable odds, we must learn from their insights and experiences along with the experiences of those who didn't make it to the top. Perhaps more important, we need to understand why existing inclusion initiatives have made so little difference. If organizations really want a representative workforce that includes more than one or two black leaders, their approach must change.

Our hope is that once companies understand the reality of the black experience, they will embrace and champion policies and programs that actually help to level the playing field—and that where there aren't yet best practices, they will begin the conversations and experiments that will lead to them. This will be hard and often uncomfortable work. But we believe it's worth it, not only for African-Americans but also for the many other underrepresented or marginalized groups. Now more than ever before, organizations and society should strive to benefit from the experiences, knowledge, and skills of all, not just a few. And while government policies can help, we believe that corporate leaders can have a much more powerful and immediate impact. As then-Senator Obama said in 2008, *"Change will not come if we wait for some other person or if we wait for some other time. We are the ones we've been waiting for. We are the change that we seek."*

The Day-to-Day Work of Diversity and Inclusion

A conversation with Airbnb's Melissa Thomas-Hunt on creating a culture in which black employees can thrive. *by Paige Cohen and Gretchen Gavett*

Most leaders of U.S. companies know that attracting diverse employees is good business. In response, the prevalence of diversity and inclusion professionals has increased and diversity trainings have become the norm. Yet these efforts, at least in their current forms, aren't boosting the representation of African-Americans in organizations and in leadership roles. What needs to change to create racially inclusive workplaces? And how can managers be the catalysts?

Dr. Melissa Thomas-Hunt is Airbnb's head of global diversity and belonging, and has been working in the diversity and inclusion space for decades, both at organizations and in academia. She spoke with us about how diversity efforts can do a better job of addressing the needs of black workers. She emphasized that there's no quick fix: "Big wins will come from interrogating seemingly mundane practices and processes, and holding managers and leaders accountable for progress toward your organization's aspirations."

An edited version of our conversation with Dr. Thomas-Hunt is below.

How do you design a diversity and inclusion program for black workers to reach leadership positions—and succeed in them?

Creating a work environment in which black employees can thrive requires deliberate, sustained efforts focused primarily in three areas: data and numbers, company culture, and day-to-day people management.

We know that numbers matter because who you hire, and at what level, directly affects the overarching narrative of what is normal and accepted in your organization. For example, if the majority of your leadership roles are occupied by white workers, you are sending the message that this group has the most potential to contribute at high levels. To change this narrative, as a first step, companies need to put more effort into increasing their pipeline of black workers.

To be clear, it's no easy feat getting black employees into an organization—and this is true globally for members of the black diaspora living in places in which they are the minority. Historical artifacts of power and privilege create all kinds of roadblocks for black people. Even when economics and levels of education are comparable, social capital—or the networks people need to gain access to opportunities—may be less available to black professionals than to their white counterparts.

So companies need to start putting more systems into place, whether through HR or recruiters, that will help them identify, attract, and hire black talent—including senior talent into critical leadership roles. But this alone is not enough. Organizations also need to make sure that the black employees they are hiring into lower-level positions are being given opportunities that set them up for success and growth. This means undergoing fundamental shifts in the cultures they create. Black employees need to enter generative work environments—ones that allow all people to grow, develop, and flourish, and ones that signal they are valued. Without these, there will be a revolving door of black talent who arrive excited, energized, and ready to contribute and leave feeling unseen and demoralized.

How can you make this culture change happen?

Though culture change is hard, and the path to it seems murky, we do know that managers are the front line. They're the ones with the power to make employees feel safe enough to contribute their knowledge and perspectives. Managers have the ability to build relationships across difference through their access to other team members and leaders. And managers can use their status to provide growth opportunities to black workers through committed sponsorship efforts and by communicating their value—including their expertise, potential, and accomplishments—to others.

Real culture change will start when managers learn how to do this, and it will require a top-down approach. Companies need to make it clear that a great supervisor is someone who creates an environment in which a diverse array of people can succeed. HR professionals need to be empowered to help managers advance inclusive behaviors and eliminate those that erode inclusion, belonging, and engagement. Resources need to be put toward training managers to understand the ways in which their own identities impact the way they engage with others. When situations arise in which black employees are experiencing microaggressions or outright discrimination, managers should know how to properly address the issues and escalate if necessary.

Organizations also need to create cultures of curiosity where people are in a constant state of discovery, learning about themselves and others. Managers can help make this happen by regularly asking their employees what they need. Holding regular check-ins with each employee is a good way for managers to demonstrate genuine interest in their team's well-being and build a foundation of trust. They should use this time to ask people if they feel supported and safe enough to contribute on a regular basis. This time will also help managers troubleshoot any issues that come up and understand their team members and aspirations, as well as how they can help them get where they want to go.

Lastly, managers should be expected to provide specific, actionable feedback to all employees and push past any hesitations they have about how that feedback will land—a fear that often stops white managers from giving black employees critical feedback. Like everyone else, black employees need honest feedback in order to grow and to get access to leadership opportunities down the line.

At some companies, talking about race consists of one formal conversation a year. How can leaders encourage more frequent discussions?

Conversations about race at work are challenging to have, or even begin, when the people involved don't have a positive relationship. That's why, at regular intervals, your employees should be encouraged to spend time with team members who appear to be different than themselves, or peers whom they do not know well. Remember that people must choose to create space for building relationships before they feel comfortable having hard conversations when racially charged situations do arise. So it's best to start building those relationships now. When opportunities for discussions surrounding race or ethnicity do come up, those participating will be more likely to assume positive intent. For the conversation to be productive, both parties need to agree that missteps will happen, and demonstrate a genuine interest in one another's experiences and perspectives.

How can you get leaders and managers on board with all of these suggestions?

Organizations need to take every opportunity to communicate what is expected of their leaders. It's not the organization's job to change attitudes. But it is their job to weave their values into the processes and practices that reinforce company culture, making sure that everyone—from individual contributors to those in leadership roles—is demonstrating behaviors that align with them.

Holding people accountable is vital to doing this successfully. If senior leaders espouse a set of values but fail to keep the people who report to them accountable for their actions or inactions, middle and lower managers will have little incentive to uphold those values and will focus instead on the business goals that are being measured. We are humans, and our attitudes are imperfect. That's why providing incentives—such as measuring diversity and inclusion efforts in performance evaluations, linking them to salary increases or other forms of compensation, and giving employees who demonstrate inclusive values public recognition—will help companies establish cultures that reinforce what they stand for. If there is a misalignment between your organization's values and the behaviors your employees exhibit, then your accountability structure is likely misaligned and needs to be rethought.

How will companies know if their diversity and inclusion programs are actually helping black employees? In what ways should they collect feedback and measure progress?

Asking employees how things are going is a good first step. To gain deeper insights, however, companies should take the bold step of analyzing employee engagement data by race and ethnicity. This is not often done because of the fear of what might be discovered, and if organizations don't look at their data by subgroup, they can easily claim that they have no knowledge of subgroup differences. Failure to measure engagement by subgroup can be perceived by black employees, and other racial or ethnic minorities, as disinterest in truly understanding the way their lived experience may diverge from others in the organization. This data should be shared and discussed internally. Where divergences in experience exist, companies must take a deep dive to understand and resolve the source of the discrepancy.

Where do you see the biggest disconnects between research and practice? And where have you seen the most promising connections?

In organizations, the degree to which we promote awareness of our unconscious biases is often held up as the solution to all the challenges that accompany diversity and inclusion efforts. But in reality, research shows that awareness can actually

increase the problematic behavior we are trying to change. This is because if we know that everyone is biased—which we are—we become less inclined to work against our own biases. We do what others do.

You've been working in inclusion for a long time. What's changed since you started? What remains stubbornly the same? And what makes you the most optimistic about the future?

More organizations recognize that they have a problem with inclusion and are committing to making changes than they were 20 years ago. I'm seeing more companies devote resources to forming diversity and inclusion programs, and hiring professionals to spearhead those efforts. There are also communities of practitioners and academics working together to identify and test best practices, whereas before, researchers and those responsible for implementing solutions rarely talked to one another, and they certainly didn't work collaboratively on challenges.

But many organizations still want quick fixes. They are impatient for better outcomes and sometimes take shortcuts. Today, certain programs still focus on "fixing" black employees as opposed to fixing organizational biases. Others showcase one-off diversity and inclusion efforts, such as showy, expensive conferences with a diverse array of speakers, yet fail to yield sustainable gains for black employees because they don't actually examine the day-to-day practices that may

undermine black employee advancement. Additionally, people remain concerned about how diversity and inclusion efforts will affect their career outcomes, and outcomes of others like themselves, if those who are historically underrepresented are given new forms of access and more developmental support.

If you had one message for other diversity and inclusion executives, what would it be?

Moving the needle on inclusion is hard. We are asking individuals to do things differently when they feel like they are already overwhelmed. Successful efforts require a deep commitment to sustained effort and offers of assistance to employees in changing their behaviors. Big wins will come from interrogating seemingly mundane practices and processes, and holding managers and leaders accountable for progress toward your organization's aspirations.

Originally published in November 2019. Reprint BG1906

The Age of Continuous Connection

by Nicolaj Siggelkow and Christian Terwiesch

A SEISMIC SHIFT IS UNDER WAY. Thanks to new technologies that enable frequent, low-friction, customized digital interactions, companies today are building much deeper ties with customers than ever before. Instead of waiting for customers to come to them, firms are addressing customers' needs the moment they arise—and sometimes even earlier. It's a win-win: Through what we call *connected strategies,* customers get a dramatically improved experience, and companies boost operational efficiencies and lower costs.

Consider the MagicBands that Disney World issues all its guests. These small wristbands, which incorporate radio-frequency-identification technology, allow visitors to enter the park, get priority access to rides, pay for food and merchandise, and unlock their hotel rooms. But the bands also help Disney locate guests anywhere in the park and then create customized experiences for them. Actors playing Disney characters, for example, can personally greet guests passing by ("Hey, Sophia! Happy seventh birthday!"). Disney can encourage people to visit

attractions with idle capacity ("Short lines at Space Mountain right now!"). Cameras on various rides can automatically take photographs of guests, which Disney can use to create personalized memory books for them, without their ever having to pose for a picture.

Similarly, instead of just selling textbooks, McGraw-Hill Education now offers customized learning experiences. As students use the company's electronic texts to read and do assignments, digital technologies track their progress and feed data to their teachers and to the company. If someone is struggling with an assignment, her teacher will find out right away, and McGraw-Hill will direct the student to a chapter or video offering helpful explanations. Nike, too, has gotten into the game. It can now connect with customers daily, through a wellness system that includes chips embedded in shoes, software that analyzes workouts, and a social network that provides advice and support. That new model has allowed the company to transform itself from a maker of athletic gear into a purveyor of health, fitness, and coaching services.

It's easy to see how Disney, McGraw-Hill, and Nike have used approaches like these to stay ahead of the competition. Many other companies are taking steps to develop their own connected strategies by investing substantially in data gathering and analytics. That's great, but a lot of them are now awash in so much data that they're overwhelmed and struggling to cope. How can managers think clearly and systematically about what

to do next? What are the best ways to use all this new information to better connect with customers?

In our research we've identified four effective connected strategies, each of which moves beyond traditional modes of customer interaction and represents a fundamentally new business model. We call them *respond to desire, curated offering, coach behavior,* and *automatic execution.* What's innovative here is not the technologies these strategies incorporate but the ways that companies deploy those technologies to develop continuous relationships with customers.

Below, we'll define these new connected strategies and explore how you can make the most of the ones you choose to adopt. But first let's take stock of the old model they're leaving behind.

Buy What We Have

Most companies still interact with customers only episodically, after customers identify their needs and seek out products or services to meet them. You might call this model *buy what we have*. In it companies work hard to provide high-quality offerings at a competitive price and base their marketing and operations on the assumption that they'll engage only fleetingly with their customers.

Idea in Brief

The Old Approach

Companies used to interact with customers only episodically, when customers came to them.

The New Approach

Today, thanks to new technologies, companies can address customers' needs the moment they arise—and sometimes even earlier. With connected strategies, firms can build deeper ties with customers and dramatically improve their experiences.

The Upshot

Companies need to make continuous connection a fundamental part of their business models. They can do so with four strategies: respond to desire, curated offering, coach behavior, and automatic execution.

Here's a typical buy-what-we-have experience: One Tuesday, working from home, David is halfway through printing a batch of urgent letters when his toner cartridge runs out. It's maddening. He *really* doesn't have time for this. Grumbling, he hunts around for his keys, gets into his car, and drives 15 minutes to the nearest office supply store. There he wanders the aisles looking for the toner section, which turns out to be an entire wall of identical-looking cartridges. After scanning the options and hoping that he recalls his printer model correctly, he finds the cartridge he needs, but only in a multipack, which is expensive. He sets off in search of a staff member who might know if the store has any single cartridges, and eventually he locates a manager, who disappears into the back of the store to check.

Much time passes. When the manager at last returns, it's to report regretfully that the store is sold out of single cartridges.

Because he has to get his letters done, David decides to buy the multipack. He grabs one and heads to the checkout counter to pay, only to find himself waiting in a long line. When he finally gets home, an hour or two later, he's not a happy guy.

We find it helpful to break the traditional customer journey into three distinct stages: *recognize,* when the customer becomes aware of a need; *request,* when he or she identifies a product or service that would satisfy this need and turns to a company to meet it; and *respond,* when the customer experiences how the company delivers the product or service. At each of these stages, David suffered a lot of discomfort, but at no point along the way did the toner company have any way of learning about his discomfort or alleviating it. Company and customer were poorly connected throughout, and both parties suffered.

It doesn't have to play out that way. Each of our four connected strategies could have helped improve David's customer experience at one or more of the stages and helped the company strengthen its business.

Let's explore specifically what each strategy entails.

Respond to Desire

This strategy involves providing customers with services and products they've requested—and doing so as quickly and seamlessly as possible. The essential capabilities here are operational: fast delivery, minimal friction, flexibility, and precise

execution. Customers who enjoy being in the driver's seat tend to like this strategy.

To provide a good respond-to-desire experience, companies need to listen carefully to what customers want and make the buying process easy. In many cases, what matters most to customers is the amount of energy they have to expend—the less, the better!

That's certainly what David wanted in his search for a toner cartridge. So let's imagine a respond-to-desire strategy that might serve him well in the future.

Say that upon realizing that he needs a replacement, David goes online to his favorite retailer, types in his printer model, and with just a click or two makes a same-day order for the correct cartridge. His credit card number and address are already stored in the system, so the whole process takes just a minute or two. A few hours later his doorbell rings, and he has exactly what he needs.

Speed is critical in a lot of respond-to-desire situations. Users of Lyft and Uber want cars to arrive promptly. Health care patients want the ability to connect at any time of day or night with their providers. Retail customers want the products they order online to arrive as quickly as possible—a desire that Amazon has famously focused on satisfying, in the process redefining how it interacts with customers. Years ago it set up a "one click" process for ordering and payment, and more recently it has gone even further than that. Today you can give Alexa a

command to order a particular product, and she'll take care of the rest of the customer journey for you. That's responding to desire.

Curated Offering

With this strategy, companies get actively involved in helping customers at an earlier stage of the customer journey: after the customers have figured out what they need but before they've decided how to fill that need. Executed properly, a curated-offering strategy not only delights customers but also generates efficiency benefits for companies, by steering customers toward products and services that firms can easily provide at the time. The key capability here is a personalized recommendation process. Customers who value advice—but still want to make the final decision—like this approach.

How might a curated-offering strategy serve David? Consider this scenario: He goes online to order his toner cartridge, and the site automatically suggests the correct one on the basis of what he has bought before. That spares him the hassle of finding the model number of his printer and figuring out which cartridge he needs. So now he just orders what the site suggests, and a few hours later, when his doorbell rings, he's had his needs smoothly and easily met.

Blue Apron and similar meal-kit providers have very effectively adopted the curated-offering strategy. This differentiates them from Instacart and many of the other grocery

delivery services that have emerged in recent years, all of which are guided by a “you order, we deliver” principle—in other words, a respond-to-desire strategy. The Instacart approach might suit you better than spending time in a supermarket checkout line, but it doesn’t relieve you of the burden of hunting for recipes and creating shopping lists of ingredients. Nor does it prevent you from overbuying when you do your shopping. Blue Apron helps on all those fronts, by presenting you with personally tailored offerings, creating an experience that many people find is more convenient, fun, and healthful than what they would choose on their own.

Coach Behavior

Both of the previous two strategies require customers to identify their needs in a timely manner, which (being human) we’re not always good at. Coach-behavior strategies help with this challenge, by proactively reminding customers of their needs and encouraging them to take steps to achieve their goals.

Coaching behavior works best with customers who know they need nudging. Some people want to get in shape but can’t stick to a workout regimen. Others need to take medications but are forgetful. In these situations a company can watch over customers and help them. Knowledge of a customer’s needs might come from information that the person has previously shared with the firm or from observing the behavior of many customers. The essential capabilities involved are a deep

understanding of customer needs ("What does the customer really want to achieve?") and the ability to gather and interpret rich contextual data ("What has the customer done or not done up to this point? Can she now enact behaviors that will get her closer to her goal?").

Here's what a coach-behavior strategy for David might look like: Perhaps the printer itself tracks the number of pages it has generated since David last changed the toner and sends that information back to the manufacturer, which knows that he will soon need a new cartridge. So it might email him a reminder to reorder. At the same time, it might encourage him to run the cleaning function on his printer—a suggestion that will help him avoid later inconveniences. Coached in this way, David will have his new printer cartridge before the old one runs out; he'll lose almost no time in replacing it; and he'll have a clean printer that performs at its best.

To implement coach-behavior approaches well, a company needs to receive information constantly from its customers so that it doesn't miss the right moment to suggest action. The technical challenge in this sort of relationship lies in enabling cheap and reliable two-way communication with customers. Traditionally, this had been difficult, but it's getting easier all the time. The advent of wearable devices, for example, allows health care companies to hover digitally over customers around the clock, constantly monitoring how they're doing.

Nike's new business model incorporates coach-behavior strategies. By making its customers part of virtual running clubs and tracking their runs, the company knows when it's time for their next workout, and through its app it can offer them audio training guides and plans. This kind of timely and personal connection builds trust and encourages customers to think of Nike as a health-and-fitness coach rather than just a shoe manufacturer, which in turn means that when the company's app nudges them to run, they're more likely to do it. This serves customers well, because it keeps them motivated and in shape. And it serves Nike well, of course, because customers who run more buy more shoes.

Automatic Execution

All the strategies we've discussed so far require customer involvement. But this last strategy allows companies to meet the needs of customers even before they've become aware of those needs.

In an automatic-execution strategy, customers authorize a company to take care of something, and from that point on the company handles everything. The essential elements here are strong trust, a rich flow of information from the customers, and the ability to use it to flawlessly anticipate what they want. The customers most open to automatic execution are comfortable having data stream constantly from their devices to companies they buy from and have faith that those companies will use their

data to fulfill their needs at a reasonable price and without compromising their privacy.

Here's how automatic execution might work for David. When he buys his printer, he authorizes the manufacturer to remotely monitor his ink level and send him new toner cartridges whenever it gets low. From then on, the onus is on the company to manage his needs, and David is spared several hassles: recognizing that he's low on toner, figuring out how to get more, and buying it. Instead, he just goes about his business. When the time is right, his doorbell will ring, and he'll have exactly what he needs.

The growing internet of things is making all sorts of automatic execution possible. David's printer cartridge scenario isn't just hypothetical: Both HP and Brother already have programs that ship replacement toner to customers whenever their printers send out a "low ink" signal. Soon our refrigerators, sensing that we're almost out of milk, will be able to order more for delivery by tomorrow morning—but naturally only after checking our calendar to make sure we're not going on a vacation and wouldn't need milk after all.

Automatic execution will make people's lives easier and in some cases will even save lives. Consider fall-detection sensors, the small medical devices worn by many seniors. Initially, the companies who made them did so using the respond-to-desire model. If an elderly person who was wearing one fell and needed help, she could press a button that activated a distress call. That

was good, but it didn't work if someone was too incapacitated to press the button. Now, though, internet-connected wearable technologies allow health care companies to monitor patients constantly in real time, which means people don't need to actively request assistance if and when they're in distress. Imagine a bracelet that monitors vital signs and uses an accelerometer to detect falls. If a person wearing the bracelet slips, tumbles down the basement stairs, and is knocked unconscious, the bracelet's sensor will immediately detect the emergency and summon help. That's automatic execution.

We're excited about automatic execution, but we want to stress that we don't see it as the best solution to all problems—or for all customers. People differ in the degree to which they feel comfortable sharing data and in having the companies serving them act on that data. One family might be delighted to receive an automatically generated personal memory book after a visit to Disney World, but another might think it's creepy and invasive. If companies want customers to make a lot of personal data available on an automated and continuous basis, they will need to prove themselves worthy of their customers' trust. They'll need to show customers that they'll safeguard the privacy and security of personal information and that they'll only recommend products and services in good faith. Breaking a customer's trust at this level could mean losing that customer—and possibly many other customers—forever.

A final important point: Given that companies are likely to have customers with different preferences, most firms will have to create a portfolio of connected strategies, which will require them to build a whole new set of capabilities. (See the table "Which connected strategies should you use?") One-size-fits-all usually won't work.

Which connected strategies should you use?

Connected strategy	Description	Key capability	Works best when	Works best for
Respond to desire	Customer expresses what she wants and when	Fast and efficient response to orders	Customers are knowledgeable	Customers who don't want to share too much data and who like to be in control
Curated offering	Firm offers tailored menu of options to customer	Making good personalized recommendations	The uncurated set of options is large and potentially overwhelming	Customers who don't mind sharing some data but want a final say
Coach behavior	Firm nudges customer to act to obtain a goal	Understanding customer needs and ability to gather and interpret rich data	Inertia and biases keep customers from achieving what's best for them	Customers who don't mind sharing personal data and getting suggestions
Automatic execution	Firm fills customer's need without being asked	Monitoring customers and translating incoming data into action	Customer behavior is very predictable and costs of mistakes are small	Customers who don't mind sharing personal data and having firms make decision for them

Repeat

Earlier, we mentioned that we like to think of the individual customer journey as having three stages: *recognize, request,* and *respond*. But there's actually a fourth stage—*repeat*—which is fundamental to any connected strategy, because it transforms stand-alone experiences into long-lasting, valuable relationships. It is in this stage that companies learn from existing interactions

and shape future ones—and discover how to create a sustainable competitive advantage.

The repeat dimension of a connected strategy helps companies with two forms of learning.

First, it allows a company to get better at matching the needs of an individual customer with the company's existing products and services. Over time and through multiple interactions, Disney sees that a customer seems to like ice cream more than fries, and theater performances more than fast rides—information that then allows the company to create a more enjoyable itinerary for him. McGraw-Hill sees that a student struggles with compound-interest calculations, which lets it direct her attention to material that covers exactly that weakness. Netflix sees that a customer likes political satire, which allows it to make pertinent movie suggestions to her.

Second, in the repeat stage companies can learn at the population level, which helps them make smart adjustments to their portfolios of products and services. If Disney sees that the general demand for frozen yogurt is rising, it can increase the number of stands in its parks that serve frozen yogurt. If McGraw-Hill sees that many students are struggling with compound-interest calculations, it can refine its online module on that topic. If Netflix observes that many viewers like political dramas, it can license or produce new series in that genre.

Both of these loops have positive feedback effects. The better the company understands a customer, the more it can customize

its offerings to her. The more delighted she is by this, the more likely she is to return to the company again, thus providing it with even more data. The more data the company has, the better it can customize its offerings. Likewise, the more new customers a company attracts through its superior customization, the better its population-level data is. The better its population data, the more it can create desirable products. The more desirable its products, the more it can attract new customers. And so on. Both learning loops build on themselves, allowing companies to keep expanding their competitive advantage.

Over time these two loops have another very important effect: They allow companies to address more-fundamental customer needs and desires. McGraw-Hill might find out that a customer wants not just to understand financial accounting but also to have a career on Wall Street. Nike might find out that a particular runner is interested not just in keeping fit but also in training to run a first marathon. That knowledge offers opportunities for companies to create an even wider range of services and to develop trusted relationships with customers that become very hard for competitors to disrupt.

We can't tell you where all this is headed, of course. But here's what we know: The age of "buy what we have" is over. If you want to achieve sustainable competitive advantage in the years ahead, connected strategies need to be a fundamental part of your business. This holds true whether you're a start-up trying to break into an existing industry or an incumbent firm trying to

defend your market, and whether you deal directly with consumers or operate in a business-to-business setting. The time to think about connected strategies is now, before others in your industry beat you to it.

Originally published in May–June 2019. Reprint R1903C

The Hard Truth about Innovative Cultures

by Gary P. Pisano

A CULTURE CONDUCIVE TO INNOVATION is not only good for a company's bottom line. It also is something that both leaders and employees value in their organizations. In seminars at companies across the globe, I have informally surveyed hundreds of managers about whether they want to work in an organization where innovative behaviors are the norm. I cannot think of a single instance when someone has said "No, I don't." Who can blame them: Innovative cultures are generally depicted as pretty fun. When I asked the same managers to describe such cultures, they readily provided a list of characteristics identical to those extolled by management books: tolerance for failure, willingness to experiment, psychological safety, highly collaborative, and nonhierarchical. And research supports the idea that these behaviors translate into better innovative performance.

But despite the fact that innovative cultures are desirable and that most leaders claim to understand what they entail, they are hard to create and sustain. This is puzzling. How can practices apparently so universally loved—even fun—be so tricky to implement?

The reason, I believe, is that innovative cultures are misunderstood. The easy-to-like behaviors that get so much attention are only one side of the coin. They must be counterbalanced by some tougher and frankly less fun behaviors. A tolerance for failure requires an intolerance for incompetence. A willingness to experiment requires rigorous discipline. Psychological safety requires comfort with brutal candor. Collaboration must be balanced with an individual accountability. And flatness requires strong leadership. Innovative cultures are paradoxical. Unless the tensions created by this paradox are carefully managed, attempts to create an innovative culture will fail.

1. Tolerance for Failure but No Tolerance for Incompetence

Given that innovation involves the exploration of uncertain and unknown terrain, it is not surprising that a tolerance for failure is an important characteristic of innovative cultures. Some of the most highly touted innovators have had their

share of failures. Remember Apple's MobileMe, Google Glass, and the Amazon Fire Phone?

And yet for all their focus on tolerance for failure, innovative organizations are intolerant of incompetence. They set exceptionally high performance standards for their people. They recruit the best talent they can. Exploring risky ideas that ultimately fail is fine, but mediocre technical skills, sloppy thinking, bad work habits, and poor management are not. People who don't meet expectations are either let go or moved into roles that better fit their abilities. Steve Jobs was notorious for firing anyone he deemed not up to the task. At Amazon, employees are ranked on a forced curve, and the bottom part of the distribution is culled. Google is known to have a very employee-friendly culture, but it's also one of the hardest places on earth to get a job (each year the company gets more than 2 million applications for about 5,000 positions). It, too, has a rigorous performance management system that moves people into new roles if they are not excelling in their existing ones. At Pixar, movie directors who cannot get projects on track are replaced.

It sounds obvious that companies should set high quality standards for their employees, but unfortunately all too many organizations fall short in this regard. Consider a pharmaceutical company I recently worked with. I learned that one of its R&D groups had not discovered a new drug candidate in more than a decade. Despite the poor performance, senior

leaders had made no real changes in the group's management or personnel. In fact, under the company's egalitarian compensation system, the scientists in the group had been receiving approximately the same salaries and bonuses as scientists in much more productive R&D units. One senior leader confided to me that short of ethics violations, the company rarely terminated anyone in R&D for subpar performance. When I asked why, he said, "Our culture is like a family. Firing people is not something we're comfortable with."

The truth is that a tolerance for failure requires having extremely competent people. Attempts to create novel technological or business models are fraught with uncertainty. You often don't know what you don't know, and you have to learn as you go. "Failures" under these circumstances provide valuable lessons about paths forward. But failure can also result from poorly thought-out designs, flawed analyses, lack of transparency, and bad management. Google can encourage risk taking and failure because it can be confident that most Google employees are very competent.

Creating a culture that simultaneously values learning through failure and outstanding performance is difficult in organizations with a history of neither. A good start is for senior leadership to articulate clearly the difference between productive and unproductive failures: Productive failures yield valuable information relative to their cost. A failure should be celebrated only if it results in learning. (The cliché "celebrating

failure" misses the point—we should be celebrating learning, not failure.) A simple prototype that fails to perform as expected because of a previously unknown technical issue is a failure worth celebrating if that new knowledge can be applied to future designs. Launching a badly engineered product after spending $500 million developing it is just an expensive flop.

Building a culture of competence requires clearly articulating expected standards of performance. If such standards are not well understood, difficult personnel decisions can seem capricious or, worse, be misconstrued as punishment for a failure. Senior leaders and managers throughout the organization should communicate expectations clearly and regularly. Hiring standards may need to be raised, even if that temporarily slows the growth of the company.

Managers are especially uncomfortable about firing or moving people when their "incompetence" is no fault of their own. Shifting technologies or business models can render a person who's very competent in one context incompetent in another. Consider how digitization has impacted the value of different skills in many industries. That sales representative whose deft interpersonal skills made him a superstar may no longer be as valuable to the organization as the introverted software engineer who develops the algorithms used to predict which customers are most likely to buy the company's products. In some cases, people can be retrained to develop new competences. But that's not always possible when really

specialized skills (say, a PhD in applied math) are needed to do a job. Keeping people who have been rendered obsolete may be compassionate, but it's dangerous for the organization.

Maintaining a healthy balance between tolerating productive failures and rooting out incompetence is not easy. A 2015 *New York Times* article about Amazon illustrates the difficulty. The piece, which was based on interviews with more than 100 current and former employees, labeled Amazon's culture as "bruising" and recounted stories of employees crying at their desks amid enormous performance pressures. One reason striking a balance is so hard is that the causes of failure are not always clear. Did a product design turn out to be flawed because of an engineer's bad judgment or because it encountered a problem that even the most talented engineer would have missed? And in the event of bad technical or business judgments, what are the appropriate consequences? Everyone makes mistakes, but at what point does forgiveness slide into permissiveness? And at what point does setting high performance standards devolve into being cruel or failing to treat employees—regardless of their performance—with respect and dignity?

2. Willingness to Experiment but Highly Disciplined

Organizations that embrace experimentation are comfortable with uncertainty and ambiguity. They do not pretend to know

all the answers up front or to be able to analyze their way to insight. They experiment to learn rather than to produce an immediately marketable product or service.

A willingness to experiment, though, does not mean working like some third-rate abstract painter who randomly throws paint at a canvas. Without discipline, almost anything can be justified as an experiment. Discipline-oriented cultures select experiments carefully on the basis of their potential learning value, and they design them rigorously to yield as much information as possible relative to the costs. They establish clear criteria at the outset for deciding whether to move forward with, modify, or kill an idea. And they face the facts generated by experiments. This may mean admitting that an initial hypothesis was wrong and that a project that once seemed promising must be killed or significantly redirected. Being more disciplined about killing losing projects makes it less risky to try new things.

A good example of a culture that combines a willingness to experiment with strict discipline is Flagship Pioneering, a Cambridge, Massachusetts, company whose business model is creating new ventures based on pioneering science. Flagship generally does not solicit business plans from independent entrepreneurs but instead uses internal teams of scientists to discover new-venture opportunities. The company has a formal exploration process whereby small teams of scientists, under the direction of one of the company's partners, undertake

research on a problem of major social or economic importance—nutrition, for example. During these explorations, teams read the literature on the topic and engage the company's broad network of external scientific advisers to conceive new scientific insights. Explorations are initially unconstrained. All ideas—however seemingly unreasonable or far-fetched—are entertained. According to founder and CEO Noubar Afeyan, "Early in our explorations, we don't ask, 'Is this true?' or 'Is there data to support this idea?' We do not look for academic papers that provide proof that something is true. Instead, we ask ourselves, 'What if this were true?' or 'If only this were true, would it be valuable?'" Out of this process, teams are expected to formulate testable venture hypotheses.

Experimentation is central to Flagship's exploration process because it is how ideas are culled, reformulated, and evolved. But experimentation at Flagship differs in fundamental ways from what I often see at other companies. First, Flagship does not run experiments to validate initial ideas. Instead, teams are expected to design "killer experiments" that maximize the probability of exposing an idea's flaws. Second, unlike many established companies that heavily fund new ventures in the mistaken belief that more resources translate into more speed and more creativity, Flagship normally designs its killer experiments to cost less than $1 million and take less than six months. Such a lean approach to testing not only enables the firm to cycle through more ideas more quickly; it also makes it

psychologically easier to walk away from projects that are going nowhere. It forces teams to focus narrowly on the most critical technical uncertainties and gives them faster feedback. The philosophy is to learn what you have gotten wrong early and then move quickly in more-promising directions.

Third, experimental data at Flagship is sacred. If an experiment yields negative data about a hypothesis, teams are expected to either kill or reformulate their ideas accordingly. In many organizations, getting an unexpected result is "bad news." Teams often feel the need to spin the data—describing the result as an aberration of some sort—to keep their programs alive. At Flagship, ignoring experimental data is unacceptable.

Finally, Flagship's venture team members themselves have a strong incentive to be disciplined about their programs. They gain no financial benefit from sticking with a loser program. In fact, just the opposite is true. Continuing to pursue a failed program means forgoing the opportunity to join a winning one. Again, compare this model with what is common in many companies: Having your program canceled is terrible news for you personally. It could mean loss of status or perhaps even your job. Keeping your program alive is good for your career. At Flagship, starting a successful venture, not keeping your program alive, is good for your career. (Disclosure: I serve on the board of a Flagship company, but the information in this example comes from a Harvard Business School case I researched and coauthored.)

Disciplined experimentation is a balancing act. As a leader, you want to encourage people to entertain "unreasonable ideas" and give them time to formulate their hypotheses. Demanding data to confirm or kill a hypothesis too quickly can squash the intellectual play that is necessary for creativity. Of course, not even the best-designed and well-executed experiments always yield black-and-white results. Scientific and business judgments are required to figure out which ideas to move forward, which to reformulate, and which to kill. But senior leaders need to model discipline by, for example, terminating projects they personally championed or demonstrating a willingness to change their minds in the face of the data from an experiment.

3. Psychologically Safe but Brutally Candid

Psychological safety is an organizational climate in which individuals feel they can speak truthfully and openly about problems without fear of reprisal. Decades of research on this concept by Harvard Business School professor Amy Edmondson indicate that psychologically safe environments not only help organizations avoid catastrophic errors but also support learning and innovation. For instance, when Edmondson, health care expert Richard Bohmer, and I conducted research on the adoption of a novel minimally invasive surgical technology by cardiac surgical teams, we

found that teams with nurses who felt safe speaking up about problems mastered the new technology faster. If people are afraid to criticize, openly challenge superiors' views, debate the ideas of others, and raise counterperspectives, innovation can be crushed.

We all love the freedom to speak our minds without fear—we all want to be heard—but psychological safety is a two-way street. If it is safe for me to criticize your ideas, it must also be safe for you to criticize mine—whether you're higher or lower in the organization than I am. Unvarnished candor is critical to innovation because it is the means by which ideas evolve and improve. Having observed or participated in numerous R&D project team meetings, project review sessions, and board of directors meetings, I can attest that comfort with candor varies dramatically. In some organizations, people are very comfortable confronting one another about their ideas, methods, and results. Criticism is sharp. People are expected to be able to defend their proposals with data or logic.

In other places, the climate is more polite. Disagreements are restrained. Words are carefully parsed. Critiques are muffled (at least in the open). To challenge too strongly is to risk looking like you're not a team player. One manager at a large company where I worked as a consultant captured the essence of the culture when she said, "Our problem is that we are an incredibly nice organization."

When it comes to innovation, the candid organization will outperform the nice one every time. The latter confuses politeness and niceness with respect. There is nothing inconsistent about being frank and respectful. In fact, I would argue that providing and accepting frank criticism is one of the hallmarks of respect. Accepting a devastating critique of your idea is possible only if you respect the opinion of the person providing that feedback.

Still, that important caveat aside, “brutally honest” organizations are not necessarily the most comfortable environments in which to work. To outsiders and newcomers, the people may appear aggressive or hard-edged. No one minces words about design philosophies, strategy, assumptions, or perceptions of the market. Everything anyone says is scrutinized (regardless of the person’s title).

Building a culture of candid debate is challenging in organizations where people tend to shy away from confrontation or where such debate is viewed as violating norms of civility. Senior leaders need to set the tone through their own behavior. They must be willing (and able) to constructively critique others’ ideas without being abrasive. One way to encourage this type of culture is for them to demand criticism of their own ideas and proposals. A good blueprint for this can be found in General Dwight D. Eisenhower’s battle-plan briefing to top officers of the Allied forces three weeks before the invasion of Normandy. As

recounted in *Eisenhower,* a biography by Geoffrey Perret, the general started the meeting by saying, "I consider it the duty of anyone who sees a flaw in this plan not to hesitate to say so. I have no sympathy with anyone, whatever his station, who will not brook criticism. We are here to get the best possible results."

Eisenhower was not just inviting criticism or asking for input. He was literally demanding it and invoking another sacred aspect of military culture: duty. How often do you demand criticism of your ideas from your direct reports?

4. Collaboration but with Individual Accountability

Well-functioning innovation systems need information, input, and significant integration of effort from a diverse array of contributors. People who work in a collaborative culture view seeking help from colleagues as natural, regardless of whether providing such help is within their colleagues' formal job descriptions. They have a sense of collective responsibility.

But too often, collaboration gets confused with consensus. And consensus is poison for rapid decision making and navigating the complex problems associated with transformational innovation. Ultimately, someone has to make a decision and be accountable for it. An accountability culture is one where individuals are expected to make decisions and own the consequences.

There is nothing inherently inconsistent about a culture that is both collaborative and accountability-focused. Committees might review decisions or teams might provide input, but at the end of the day, specific individuals are charged with making critical design choices—deciding which features go and stay, which suppliers to use, which channel strategy makes most sense, which marketing plan is best, and so on. Pixar has created several ways to provide feedback to its movie directors, but as Ed Catmull, its cofounder and president, describes in his book *Creativity, Inc.*, the director chooses which feedback to take and which to ignore and is held accountable for the contents of the movie.

Accountability and collaboration can be complementary, and accountability can drive collaboration. Consider an organization where you personally will be held accountable for specific decisions. There is no hiding. You own the decisions you make, for better or worse. The last thing you would do is shut yourself off from feedback or from enlisting the cooperation and collaboration of people inside and outside the organization who can help you.

A good example of how accountability can drive collaborative behavior is Amazon. In researching a case for Harvard Business School, I learned that when Andy Jassy became head of Amazon's then-fledgling cloud computer business, in 2003, his biggest challenge was figuring out what services to build (hardly an easy task given that cloud services

were a completely new space for Amazon—and the world). Jassy immediately sought help from Amazon's technology teams, its business and technical leaders, and external developers. Their feedback about requirements, problems, and needs was critical to the early success of what eventually became Amazon Web Services—today a profitable $12 billion business run by Jassy. For Jassy, collaboration was essential to the success of a program for which he was personally accountable.

Leaders can encourage accountability by publicly holding themselves accountable, even when that creates personal risks. Some years ago, when Paul Stoffels headed R&D at Johnson & Johnson's pharmaceutical division, his group experienced a failure in a major late-stage clinical program. (Disclosure: I have consulted for various divisions of Johnson & Johnson.) As Stoffels recounted at a meeting of J&J managers that I attended, senior leadership and the board demanded to know who was at fault when the program had its setback. "I am accountable," Stoffels replied. "If I let this go beyond me, and I point to people who took the risk to start and manage the program, then we create a risk-averse organization and are worse off. This stops with me." Stoffels, now chief scientific officer for J&J, shares this story frequently with employees throughout the corporation. He finishes with a simple promise: "You take the risk; I will take the blame." And then he urges his audience to cascade this principle down the organization.

5. Flat but Strong Leadership

An organizational chart gives you a pretty good idea of the structural flatness of a company but reveals little about its cultural flatness—how people behave and interact regardless of official position. In culturally flat organizations, people are given wide latitude to take actions, make decisions, and voice their opinions. Deference is granted on the basis of competence, not title. Culturally flat organizations can typically respond more quickly to rapidly changing circumstances because decision making is decentralized and closer to the sources of relevant information. They tend to generate a richer diversity of ideas than hierarchical ones, because they tap the knowledge, expertise, and perspectives of a broader community of contributors.

Lack of hierarchy, though, does not mean lack of leadership. Paradoxically, flat organizations require stronger leadership than hierarchical ones. Flat organizations often devolve into chaos when leadership fails to set clear strategic priorities and directions. Amazon and Google are very flat organizations in which decision making and accountability are pushed down and employees at all levels enjoy a high degree of autonomy to pursue innovative ideas. Yet both companies have incredibly strong and visionary leaders who communicate goals and articulate key principles about how their respective organizations should operate.

Here again, the balance between flatness and strong leadership requires a deft hand by management. Flatness does not mean that senior leaders distance themselves from operational details or projects. In fact, flatness allows leaders to be closer to the action. The late Sergio Marchionne, who led the resurrection of first Fiat and then Chrysler (and was the architect of their merger), commented to me during an interview for a Harvard Business School case I wrote: "At both companies, I used the same core principles for the turnaround. First, I flattened the organization. I had to reduce the distance between me and the people making decisions. [At one point, Marchionne had 46 direct reports between the two organizations.] If there is a problem, I want to know directly from the person involved, not their boss."

At both Fiat and Chrysler, Marchionne moved his office to the engineering floor so that he could be closer to product planning and development programs. He was famous both for being detail oriented and for pushing decision making down to lower levels in the organization. (With so many direct reports, it was nearly impossible for him not to!)

Getting the balance right between flatness and strong leadership is hard on top management and on employees throughout the organization. For senior leaders, it requires the capacity to articulate compelling visions and strategies (big-picture stuff) while simultaneously being adept and competent with technical and operational issues. Steve Jobs was a great

example of a leader with this capacity. He laid out strong visions for Apple while being maniacally focused on technical and design issues. For employees, flatness requires them to develop their own strong leadership capacities and be comfortable with taking action and being accountable for their decisions.

Leading the Journey

All cultural changes are difficult. Organizational cultures are like social contracts specifying the rules of membership. When leaders set out to change the culture of an organization, they are in a sense breaking a social contract. It should not be surprising, then, that many people inside an organization—particularly those thriving under the existing rules—resist.

Leading the journey of building and sustaining an innovative culture is particularly difficult, for three reasons. First, because innovative cultures require a combination of seemingly contradictory behaviors, they risk creating confusion. A major project fails. Should we celebrate? Should the leader of that program be held accountable? The answer to these questions depends on the circumstances. Was the failure preventable? Were issues known in advance that could have led to different choices? Were team members transparent? Was there valuable learning from the experience? And so on. Without clarity

around these nuances, people can easily get confused and even cynical about leadership's intentions.

Second, while certain behaviors required for innovative cultures are relatively easy to embrace, others will be less palatable for some in the organization. Those who think of innovation as a free-for-all will see discipline as an unnecessary constraint on their creativity; those who take comfort in the anonymity of consensus won't welcome a shift toward personal accountability. Some people will adapt readily to the new rules—a few may even surprise you—but others will not thrive.

Third, because innovative cultures are systems of interdependent behaviors, they cannot be implemented in a piecemeal fashion. Think about how the behaviors complement and reinforce one another. Highly competent people will be more comfortable with decision making and accountability—and their "failures" are likely to yield learning rather than waste. Disciplined experimentation will cost less and yield more useful information—so, again, tolerance for failed experiments becomes prudent rather than shortsighted. Accountability makes it much easier to be flat—and flat organizations create a rapid flow of information, which leads to faster, smarter decision making.

Beyond the usual things that leaders can do to drive cultural change (articulate and communicate values, model target behaviors, and so on), building an innovative culture requires some specific actions. First, leaders must be very transparent

with the organization about the harder realities of innovative cultures. These cultures are not all fun and games. Many people will be excited about the prospects of having more freedom to experiment, fail, collaborate, speak up, and make decisions. But they also have to recognize that with these freedoms come some tough responsibilities. It's better to be up-front from the outset than to risk fomenting cynicism later when the rules appear to change midstream.

Second, leaders must recognize that there are no shortcuts in building an innovative culture. Too many leaders think that by breaking the organization into smaller units or creating autonomous "skunk works" they can emulate an innovative start-up culture. This approach rarely works. It confuses scale with culture. Simply breaking a big bureaucratic organization into smaller units does not magically endow them with entrepreneurial spirit. Without strong management efforts to shape values, norms, and behaviors, these offspring units tend to inherit the culture of the parent organization that spawned them. This does not mean that autonomous units or teams can't be used to experiment with a culture or to incubate a new one. They can. But the challenge of building innovative cultures inside these units should not be underestimated. And they will not be for everyone, so you will need to select very carefully who from the parent organization joins them.

Finally, because innovative cultures can be unstable, and tension between the counterbalancing forces can easily be

thrown out of whack, leaders need to be vigilant for signs of excess in any area and intervene to restore balance when necessary. Unbridled, a tolerance for failure can encourage slack thinking and excuse making, but too much intolerance for incompetence can create fear of risk taking. Neither of these extremes is helpful. If taken too far, a willingness to experiment can become permission to take poorly conceived risks, and overly strict discipline can squash good but ill-formed ideas. Collaboration taken too far can bog down decision making, but excessive emphasis on individual accountability can lead to a dysfunctional climate in which everyone jealously protects his or her own interests. There is a difference between being candid and just plain nasty. Leaders need to be on the lookout for excessive tendencies, particularly in themselves. If you want your organization to strike the delicate balance required, then you as a leader must demonstrate the ability to strike that balance yourself.

Originally published in January–February 2019. Reprint R1901C

Creating a Trans-Inclusive Workplace

by Christian N. Thoroughgood, Katina B. Sawyer, and Jennica R. Webster

FOR MOST OF US, work is stressful in and of itself. Imagine carrying the added emotional weight of having to deny and suppress one of the most fundamental aspects of who you are—your gender identity—because it doesn't conform with society's norms regarding gender expression. And imagine how it would feel if you revealed your authentic self to those you work with and see every day, only to have them reject, ostracize, or ignore you as a result. (Maybe you do not have to imagine at all.)

These issues are pervasive for many trans people, who often experience stigma and discrimination, hostility, and pressure to "manage" their identities in social settings—including the workplace—to suit the expectations of others. Such experiences can set in motion a host of psychological responses that have devastating consequences for trans individuals' emotional

well-being, job satisfaction, and inclination to remain with an employer.

Despite a growing global awareness of the struggles trans people face, many employers remain ill-equipped to create the policies and workplace cultures that would support trans employees. Part of the problem is a lack of knowledge about these challenges. Indeed, even companies that are LGBTQ+-friendly usually focus more on the “LGB” than on the “TQ+.”

The overriding reason to address this issue is that it’s simply the right thing to do. Nobody who works hard and contributes to an organization’s success should ever have to feel stigmatized and fearful of coming to work each day. But that’s not the only reason. A failure to adopt trans-specific policies and practices can cost businesses dearly in the form of higher turnover, decreased engagement and productivity, and possible litigation. Discriminatory behavior in general also hurts the company’s brand.

Fortunately, research on how employers can more effectively attract, retain, and promote the well-being and success of their trans employees is growing. Although we are not members of the trans community, we’ve spent the past seven years learning from a diverse population of trans people in the course of our research as organizational psychologists specializing in gender-related issues. We’ve interviewed and surveyed more than 1,000 trans employees from a range of industries and

professions throughout North America. In this article we share their voices and experiences and outline what we've learned.

The Roots of Stigma and Discrimination

Why do trans individuals so often face stigma and discrimination? The answer resides in how people are socialized to understand and enact gender. A large body of scholarly research in social and developmental psychology has demonstrated that gendered behavior is *learned*: From a young age, boys and girls are encouraged to display stereotypically gendered behaviors and discouraged from displaying non-normative ones. Just think about the tradition of giving pink items to baby girls and blue items to baby boys. The preference for these colors has no biological roots; in fact, pink was once considered the more "masculine" color. Yet over time little boys come to prefer blue and little girls come to prefer pink; they are subtly rewarded for liking their respective colors and may even be chastised for liking the other color. Moreover, children pick up on subtle signals from their parents and important others who enforce gender stereotypes. For example, when donning female garments during dress-up, girls might be told they look pretty, while boys might be told they look silly. Children seek to fulfill gender expectations in order to secure parental and, later, peer acceptance. As we grow up, it becomes difficult to distinguish between expressions of gender

we actually prefer and those we have been socially rewarded for.

Idea in Brief

The Problem

Trans people often experience stigma and discrimination, hostility, and pressure to “manage” their identities in social settings, including the workplace.

Why It Occurs

Despite a growing global awareness of these struggles, many employers remain ill- equipped to develop policies and workplace cultures that support trans employees.

What to Do

Research and interviews or surveys of more than 1,000 trans people suggest four things companies can do: adopt basic practices of trans inclusivity involving bathroom use, dress codes, and pronouns; support gender transitions; develop trans-specific diversity trainings; and utilize resiliency interventions.

As a result of this socialization, gender norms provide perhaps the most basic organizing framework by which people define themselves and others. And because they are widely shared and deeply rooted, they are extremely difficult to change. Thus trans people face a unique quandary. For example, when a trans woman—whose sex was assigned male at birth and who knows herself to be female—adopts typically

female clothing and jewelry, she breaks with expectations regarding how she should define and express her gender.

Unfortunately, such situations most often mean that trans individuals are stigmatized—that is, socially devalued—providing a basis for discrimination against them. Studies suggest that the costs of that stigma and discrimination are steep. For example, a 2015 survey of 27,715 trans individuals residing in the United States revealed that a staggering 77% of those who had held a job in the year prior took active steps to avoid mistreatment at work, such as hiding their gender identity, delaying their gender transition (or living as their true selves only after work and on weekends), refraining from asking their employers to use their correct pronouns (*he, she, they, ze*), or quitting their jobs. Sixty-seven percent reported negative outcomes such as being fired or forced to resign, not being hired, or being denied a promotion. And nearly a quarter reported other types of mistreatment based on their gender identity or expression—for example, being required to present as the sex assigned to them at birth to keep a job, having private information about their trans identity shared without permission, or being denied access to bathrooms that align with their gender identity. Such experiences may be compounded for a trans person who holds more than one stigmatized identity—for example, a black trans woman.

Research also suggests that stigma and discrimination can result in ruminative thoughts, a negative self-image,

hopelessness, social isolation, and alcohol abuse or other dysfunctional coping behaviors. Such responses pave the way for even greater mental health challenges, including major depression and anxiety.

In one of our own investigations, we collected daily survey data from 105 trans employees in the United States across two workweeks. The results revealed that 47% of participants experienced at least some discriminatory behavior on a daily basis at work, such as being the target of transphobic remarks, being ignored, or being pressured to act in “traditionally gendered” ways. They reported robust increases in hypervigilance and rumination at work the day following such an experience. The extent to which they had to be “on guard” around their coworkers and try to make sense of negative events predicted their emotional exhaustion during the workday.

In another study, this one involving 165 trans employees from various industries and occupations in North America, we replicated those results and extended them to other outcomes, including diminished job satisfaction and a greater desire to quit. One trans woman, an educator, who felt deeply unsupported by the administration after she reported being harassed, told us, “Students were being removed from my class, rumors were spread about me, and it just wasn’t a great place to be working anymore.” Another trans woman, who worked in retail, recalled that her direct supervisor joked about

trans individuals and that customers would tell her not to bring her "lifestyle" into the workplace. As a result, she said, "I'm constantly aware of who is around me at all times. And when I'm around other people, it makes me very unsettled." A trans man in the business sector echoed this intense sense of distress: "Most of my stress that comes from work is related to just anxiety and worry [about interactions with coworkers], just constantly wondering about things that have happened and what might happen."

Employers should be aware of the business costs of ignoring these issues. A March 2012 report by the Center for American Progress indicated that companies in the United States lose an estimated $64 billion annually as a result of having to replace employees who departed because of unfairness and discrimination; many of those individuals were members of the LGBTQ+ community.

Hostility and discrimination also increase absenteeism, undermine commitment and motivation, and decrease productivity. A recent study by the Human Rights Campaign found that employee engagement declines by as much as 30% in unfriendly work environments. Although the study focused on LGBTQ+ employees more broadly, its findings are no doubt representative of trans people's experiences. In addition to hiding who they are at work, which LGB individuals often must do with respect to their sexual identity, trans people must hide

their gender expression, including how they dress, speak, and present themselves.

Discriminatory workplaces also prevent companies from attracting and retaining top talent. When employers, whether knowingly or unknowingly, fail to address prejudicial behavior, they send a potent message about their indifference and develop an external reputation for being an unwelcoming place to work. (According to the Level Playing Field Institute, one in four people who experience unfairness in the workplace report being highly unlikely to recommend their organization to others.) Furthermore, laws relating to gender identity and expression, although still severely lacking in the aggregate, are evolving at the local, state, and federal levels—creating greater obligations for employers. Without comprehensive strategies for addressing issues around gender identity and expression, organizations risk being sued. Those legal actions can be expensive to litigate, distracting to business activities, and damaging to a company's reputation, in addition to involving costly payouts. But it is our hope that companies will approach trans inclusivity from a moral and ethical standpoint rather than a purely economic one.

Supporting Your Trans Workforce

Organizations should not wait for the courts to determine that trans individuals are fully protected under the law. (See the

sidebar “Gender Expression and Employment Law.”) Instead they should proactively incorporate gender-identity-specific nondiscrimination policies and practices throughout their businesses. That involves two key issues: protecting and promoting the rights of people of all gender identities and expressions, and increasing employees’ understanding and acceptance of their trans colleagues. In a meta-analysis we conducted with Cheryl Maranto and Gary Adams, we found strong links between the degree to which employers enact these practices and the job attitudes, psychological well-being, and disclosure decisions of LGBTQ+ community members. In another study, focused specifically on trans employees, Enrica Ruggs and her coauthors found that the presence of trans-supportive policies was positively related to participants’ openness about their identities and their decreased experiences of discrimination at work. However, such effects are likely to occur only when leaders model these policies consistently in both words and behavior. Also, it should be noted that effective diversity and equity practices have been found to positively impact the productivity of all employees.

Gender Expression and Employment Law

LAWS REGARDING GENDER and gender expression are constantly evolving and differ according to location. In the United States no federal law prohibits discrimination against trans people, and only 19

states have explicit protections for trans workers. Additionally, the Religious Freedom Restoration Act of 1993 makes it more difficult for trans employees to file discrimination complaints against employers who justify their practices on religious grounds. Using religious freedom as a rationale, certain states have enacted laws to revoke or prohibit equal protections for trans individuals. Although gender expression has been covered in some court cases under the broader sex-discrimination protections within Title VII of the Civil Rights Act, in the absence of a federal law it remains up to the courts to decide case outcomes according to their interpretations of prior case law. Indeed, the U.S. Supreme Court in 2019 began deliberating over whether Title VII sex protections extend to LGBTQ+ populations.

At the global level, laws regarding gender expression vary widely. Many countries, including the United Kingdom, Spain, and South Africa, have trans-specific antidiscrimination protections. However, being trans is punishable by law in countries such as Saudi Arabia, Nigeria, and Malaysia. In many other countries, as in the United States, being trans is neither punishable nor protected, leaving oft-discriminated-against trans people in a state of uncertainty regarding their status as equal citizens under the law. When doing business in a global environment, it is vital to be mindful of how protections may vary and what this may mean for the safety and well-being of trans employees. Even when operating within intolerant cultural contexts, it is important to practice inclusivity consistently.

Here are four practices that we recommend employers adopt. Further resources can be found through professional associations such as the Society for Human Resource Management and nonprofit organizations such as the Human Rights Campaign, Out & Equal, and the Transgender Law Center.

1. Adopt basic trans-inclusive policies

An extensive body of social psychology research suggests that human beings are highly attuned to signals regarding the value ascribed to them by others. To one degree or another, we all have a basic need to belong and a prewired, unconscious monitoring system that tracks the quality of our relationships. When we detect signs of social devaluation (apathy, disapproval, or rejection), we experience negative emotions and a loss of self-esteem. When we detect signs of social valuation (praise, affection, or admission to a desired group), just the opposite occurs. Thus inclusive policies and practices—such as those related to bathroom access, dress codes, and pronoun and name usage—send vital messages to trans employees about their value as organizational members.

Bathroom access. Instituting gender-neutral bathrooms or encouraging trans employees to use bathrooms that align with their gender identity is one important way to signal to those employees that they are valued. Diversity trainings should educate other employees on the importance of being accepting and welcoming when they find themselves in a company bathroom with a trans coworker. One of our participants, a trans man working in business, said, "When I started using the men's room at work, a number of men didn't like it. An engineer, a cisgender man in his forties who didn't work with

me directly, went out of his way to make me feel safe and welcome in the men's room, and I was extremely grateful."

Some have suggested that allowing employees to use bathrooms that align with their gender identity will increase the risk of sexual harassment and assault against women. But a 2018 report published in *Sexuality Research and Social Policy* suggests that such incidents in bathrooms are rare, regardless of any gender-identity policy on bathroom usage. In fact, harassment and assault generally are most often perpetrated by straight, cisgender males against straight, cisgender females.

Dress codes. Some organizations, including Accenture, have begun to regionally implement gender-neutral dress codes. By making explicit that all employees may select from a range of options, such as dress shirts, pantsuits, and skirt suits, companies can help destigmatize varying expressions of gender. Such policies may also aid in recruitment and retention by signaling that normativity is not expected.

Pronoun and name usage. Another way to signal to trans employees that they are valued is to pay serious attention to their correct names and pronouns. Many trans people identify on the traditional binary scale—as either male or female—and thus use *he, him,* and *his* or *she, her,* and *hers* as pronouns. Yet many others who also fall under the broad category "trans"—such as genderqueer, gender-fluid, and nonbinary individuals

—use alternative pronouns, such as *they, them,* and *theirs* or *ze, zir,* and *zem*.

It’s clear from our conversations and research that the “misgendering” of trans employees, whether intentional or unintentional, is relatively common at work. A onetime slipup—such as using an incorrect pronoun for a colleague who has recently transitioned—may be considered an honest mistake. (One should apologize, move on, and make sure to get it right the next time.) Using the right pronouns and names on a regular basis can be more meaningful than one might think. When asked to reflect on courageous acts coworkers had performed in support of the rights of trans employees, many of our participants recalled instances in which a cisgender employee guided others on proper pronoun usage. A simple “Katie uses ‘she’ as a pronoun” works, as does a gentle correction: “Have you seen him?” “Yes, I saw her in the conference room.”

Employers can address this issue in several ways. First, they can keep records of employees’ chosen names and correct pronouns; this helps ensure that whenever possible, appropriate terms will be used for personnel and administrative purposes, such as directories, email addresses, and business cards. Second, encourage all employees to use name badges and email signatures that include their desired names and correct pronouns; this enables people to learn those names and pronouns and cultivates awareness of the varying gender

identities that colleagues may possess. Third, take advantage of training programs, onboarding initiatives, and employee handbook content to make clear that proper pronoun usage is part of creating an environment in which all employees feel valued and respected. Goldman Sachs, for example, recently launched an internal campaign to make employees more aware of the importance of pronouns and to encourage them to proactively share their pronouns with colleagues.

2. Support gender transitions

Transitioning is not a single event but, rather, a *process,* which begins with a deeply personal decision that usually results from years of soul-searching. The decision to come out, or disclose, at work is also complicated. People weigh the positive consequences of doing so (freedom from living a "double life" and expression of one's true self) against the negative ones (potential rejection and career ramifications). One of our study participants, a trans woman in the transportation industry, told us, "After nearly a year of soul-searching, research, therapy, support group attendance, and deep personal reflection, I 'came out' to my supervisor as transgender. … I finished talking, paused, and waited for her reply. My heart was in my throat. I knew this meeting might forever change the way she thought of me, and that I could not un-say what had been said."

Then the woman recounted her boss's reaction: "After a few moments, her very first words were 'We're not just a team here, we're a family, and this is your home. You have the right to be who you are and to be treated with respect and dignity. I will do everything I can to make sure your transition is as smooth and trouble-free as it can be.' She then got busy arranging meetings with the head of the department and the head of HR."

Someone deciding to transition chooses what that process will look like and how long it will take. A transition may involve gender-confirmation surgery (not all trans people undergo medical procedures). Some gender-fluid individuals spend their lives transitioning between and within various gender expressions, as they continually reinterpret and redefine themselves. Employers must develop a comprehensive approach to managing gender transitions—one that focuses on the employee but also on cultivating a work environment conducive to the transition process.

First, helping transitioning employees who elect medical procedures to cover costs—and making sure they have access to health care benefits that are gender-identity-specific—can reduce the stress and anxiety of coming out at work. Such commitment sends a highly affirming message to trans employees about their value.

Second, it is paramount that employees be asked what they need during their transitions and how they would like the

process handled. Only by listening to and collaborating with them can employers ensure that people are not inadvertently "outed" without permission or before they're ready.

Gender Identity and Expression: A Glossary

PEOPLE HAVE DIFFERING LANGUAGE to describe who they are and how they want to label their identities. The terms below are frequently used, but we acknowledge that these and other definitions are constantly evolving. Further, it's important to note that individuals know their own identity best and should always be consulted about how they'd like to be referred to. (For more, see the Human Rights Campaign's glossary of terms.)

Cisgender: A gender identity that aligns with the sex assigned at birth.

Gender expression: The ways in which people—trans or not—choose to convey their gender identity through dress, verbal communication styles, and other outward behavior.

Gender fluid: Refers to people who feel more male, more female, or some combination of the two at various times, and who therefore express their gender identity more dynamically over time.

Gender identity: How one understands one's own gender, regardless of the sex assigned at birth.

Genderqueer: A gender identity and expression that are not tied to a traditional male/female view of the gender spectrum. Those who identify as genderqueer may identify as men or women, as neither, or as some combination of the two.

Trans: An umbrella term for cases in which gender does not align with societal expectations regarding the sex assigned at birth. Some people who fall under the umbrella decide to transition; others do not, because they don't define themselves according to the traditional male-female binary or because they have a more fluid view of their identity over time.

Transgender: A gender identity that does not align with the sex assigned at birth. For example, a transgender woman is someone whose sex assigned at birth was male.

Third, if approached by an employee, an HR manager can provide information concerning where to learn more about treatment options, organizational support groups, and other available resources and can develop strategies to help the employee manage work/life issues that may arise during the process. Including direct supervisors in such meetings, if the employee feels comfortable with this, can promote empathy and aid in crafting flexible and informed plans adapted to each individual's unique needs. Google, Cigna, and Chevron have implemented such initiatives.

Fourth, and equally important, our research suggests that leaders and managers must proactively cultivate a supportive work environment. The period of transitioning is particularly sensitive; indeed, individuals may be ostracized or pressured by peers to suppress their identities during this time, increasing their susceptibility to depression, anxiety, and even suicidal thinking. Moreover, any trans person seeking surgery will be questioned by the surgical team about the existence of support networks, which are often required for someone who is seeking gender-confirming procedures. Thus having supportive policies and plans in place will remove one or more barriers to care for trans employees.

Authority figures who model trans-inclusive behaviors on a consistent basis are crucial to creating a supportive environment. Many of our participants said they would not have felt comfortable inquiring about transition benefits, much less been successful in their transitions, if senior leaders and frontline managers had not shown support, which tends to have a trickle-down effect on lower-level employees. Top leaders can do this in various ways, such as by attending or presenting at conferences about trans-specific issues, publicly championing gender-inclusive dress codes and bathroom usage initiatives, and using their correct names and gender pronouns.

Of course coworkers play a key role as well. In a recent study using interview and survey data from 389 trans employees and conducted with Larry Martinez, Enrica Ruggs, and Nicholas

Smith, we found that those who were relatively far along in their transitions were more satisfied with their jobs, felt a greater sense of "fit" in their workplaces, and reported less discrimination than those who had not transitioned or were less far along in the process. We also found that this effect was explained *not* by participants' sense of consistency between their inner gender identity and their outward expression of gender—what is referred to as *action authenticity*—but, rather, by the perception that coworkers had the same understanding of their gender that the participants did, which is known as *relational authenticity*. One participant, a trans man who works as a museum curator, said, "There was a point where people started seeing me as just one of the guys. And I think that at that point I started feeling like I fit in a lot better. It's the individuals [coworkers] who make that possible." In a poignant example from a separate study, a trans woman in manufacturing reported a moment at a company function: "I appeared in a dress for the first time at a party. One of the housekeeping aides grabbed my hand and pulled me onto the dance floor in front of everyone. His courage in accepting who I was in front of all our coworkers can bring me to tears to this day."

To help in cultivating supportive relationships, work groups should be told when those who are transitioning will be out of the office, whether they will return part-time, and what work will have to be covered during their absence. Emphasizing the

need for coworkers to show sensitivity, provide emotional support, and act in ways that affirm the gender identity of their colleagues is crucial. For example, people can make it clear that they are available to talk about any issues related to transitioning or gender expression—while following trans employees' lead about when and where to have those conversations. That approach enhances feelings of support and care and allows trans employees to be comfortable having honest conversations with their colleagues. Even well-intentioned employees may be nervous about their ability to support a colleague through a transition, and employers can help ease some of their anxiety by taking the above steps.

3. Develop trans-specific diversity training

More general training on gender-identity topics is also essential. Although media coverage has helped facilitate conversations about gender identity and expression, corporate diversity trainings still have room for improvement. We offer two recommendations:

1. **Include contact with those who identify along the trans identity spectrum.** A large body of research on the "contact hypothesis" suggests that providing opportunities to build relationships with specific groups—to hear their stories, appreciate their challenges, and gain empathy—is critical for shifting attitudes and behavior toward them. However, it is not the responsibility

of members of the LGBTQ+ community to educate others or to be visible in this way; “out” trans employees should be included in trainings only if they are willing. If they’re not, many corporate training firms and LGBTQ+ nonprofit organizations offer training of this nature.

2. **Help cisgender employees develop the skills to become informal champions of their transgender colleagues.** Research suggests that many people lack the knowledge and confidence to challenge prejudice. That’s why some companies have sought to equip their employees, especially leaders, with concrete strategies for stepping out of their comfort zones and engaging in “courageous conversations” regarding difficult diversity-related topics. For example, an employee who witnesses biased behavior is encouraged to respectfully but directly call it out. That might mean pulling someone aside to explain the potential damage from a biased comment, or having coffee with someone to tactfully share why a behavior was noninclusive. The chairman of PwC launched the CEO Action for Diversity & Inclusion coalition to normalize diversity-related conversations across top-level leaders in large companies. At Bank of America employees are encouraged to discuss gender, race, and other identity-related issues in a respectful, learning-focused manner.

These efforts pay off. In a forthcoming study we will report that cisgender employees who challenge noninclusive policies

and behavior send an important message of inclusion to their trans colleagues. Our findings suggest that these behaviors may come in three related forms: *advocacy,* such as taking the initiative to publicly support trans causes; *defending,* such as protecting trans coworkers from judgment or hostility; and *educating,* such as spreading awareness of trans issues in the organization. We found that trans individuals who had recently witnessed these behaviors tended to report an increased sense of worth as organizational members, were more satisfied with their jobs, and were less emotionally depleted by work.

One trans man in government recalled feeling immense gratitude toward his assistant when she spoke out after he was treated poorly by a manager. “This came about as I sat at a lunch table at an empty chair,” he recalled. “When he saw I was sitting there, [he] jumped up like he had sat next to a very large spider. She [my assistant] voiced, ‘Scott, that was so rude’—twice! That brought me to an island of relief.” Courageous acts like this predicted individuals’ job satisfaction and well-being a full six weeks later.

Despite the good intentions of many cisgender employees, however, trans people may not always want others to represent their interests, especially when those others lack in-depth knowledge of the various issues, challenges, and nuances surrounding their work and life experiences. And research suggests that employees who possess a “savior mentality” (that is, are motivated by a desire to be perceived as good people)

may end up doing more harm than good. Accordingly, HR practitioners should train employees to appropriately ask whether trans colleagues prefer to speak up for themselves. (If they wish to be, trans employees should be involved in this training.) The simple act of consulting before taking action gives a trans person agency and autonomy in deciding how the situation should be handled.

4. Utilize interventions to build resiliency

Research also supports the idea that trans individuals can benefit from interventions to help them manage their stress. In a recent two-week experience sampling study of ours, we found evidence to suggest that mindfulness—a state of nonjudgmental attention to present-moment experiences—can insulate trans employees from emotional exhaustion the day after experiencing a stigmatizing event at work. This effect was explained by a reduction in defensive, distrustful patterns of thinking such as hypervigilance and rumination.

Unfortunately, it's not realistic to assume that prejudice toward trans employees will be eliminated quickly and easily through workplace initiatives. Such changes take time. And although the main goal of employers should be to root out prejudice at a structural level through formal diversity policies and practices, it's also important to offer tools—such as mindfulness training, cognitive behavioral training, and self-

compassion training—for reducing the harmful outcomes that stigma creates in marginalized populations.

Only when people feel totally authentic and connected with their organizations can they achieve their full potential at work. Trans employees are no exception. Yet few companies have succeeded in creating an inclusive work environment for people who don't identify with societal gender norms. We hope that the research and the proactive steps we've outlined will help change that. Employers that get this right aren't just being savvy from a business standpoint. They are also crafting a corporate legacy—one in which human dignity is prioritized and doing the right thing by employees is regarded as fundamental to success.

Originally published in March–April 2020. Reprint R2002J

When Data Creates Competitive Advantage

… and When It Doesn’t

by Andrei Hagiu and Julian Wright

MANY EXECUTIVES AND INVESTORS ASSUME that it’s possible to use customer-data capabilities to gain an unbeatable competitive edge. The more customers you have, the more data you can gather, and that data, when analyzed with machine-learning tools, allows you to offer a better product that attracts more customers. You can then collect even more data and eventually marginalize your competitors in the same way that businesses with sizable network effects do. Or so the thinking goes. More often than not, this assumption is wrong. In most instances people grossly overestimate the advantage that data confers.

The virtuous cycles generated by data-enabled learning may look similar to those of regular network effects, wherein an offering—like a social media platform—becomes more valuable as more people use it and ultimately garners a critical mass of users that shuts out competitors. But in practice regular

network effects last longer and tend to be more powerful. To establish the strongest competitive position, you need them *and* data-enabled learning. However, few companies are able to develop both. Nevertheless under the right conditions customer-generated data can help you build competitive defenses, even if network effects aren't present. In this article we'll walk you through what those conditions are and explain how to evaluate whether they apply to your business.

What Has Changed?

Companies built on data have been around for a long time. Take credit bureaus and the information aggregators LexisNexis, Thomson Reuters, and Bloomberg, just to name a few. Those companies are protected by significant barriers to entry because of the economies of scale involved in acquiring and structuring huge amounts of data, but their business models don't involve gleaning data from customers and mining it to understand how to improve offerings.

Gathering customer information and using it to make better products and services is an age-old strategy, but the process used to be slow, limited in scope, and difficult to scale up. For automakers, consumer-packaged-goods companies, and many other traditional manufacturers, it required crunching sales data, conducting customer surveys, and holding focus groups. But the sales data often wasn't linked to individual customers,

and since surveys and focus groups were expensive and time-consuming, only data from a relatively small number of customers was collected.

That changed dramatically with the advent of the cloud and new technologies that allow firms to quickly process and make sense of vast amounts of data. Internet-connected products and services can now directly collect information on customers, including their personal details, search behavior, choices of content, communications, social media posts, GPS location, and usage patterns. After machine-learning algorithms analyze this "digital exhaust," a company's offerings can be automatically adjusted to reflect the findings and even tailored to individuals.

These developments make data-enabled learning much more powerful than the customer insights companies produced in the past. They do not, however, guarantee defensible barriers.

Idea in Brief

The Assumption

Companies can build winner-take-all positions by collecting and analyzing customer data. The more customers a firm has, the more data it can gather and mine; the resulting insights allow it to offer a better product that attracts even more customers, from which it can collect still more data.

The Reality

Even when customer data does confer a competitive advantage, it gives rise to network effects only infrequently. And that advantage

may not last.

The Solution

To understand the edge that data-enabled learning can provide, companies should answer seven questions, which examine the value of the data; whether its marginal value drops quickly; how fast it becomes obsolete; whether it's proprietary; whether the improvements from it can be easily imitated; whether they enhance the product for current users, other users, or both; and how fast insights can be incorporated into products.

Building Moats with Data-Enabled Learning

To determine to what degree a competitive advantage provided by data-enabled learning is sustainable, companies should answer seven questions:

1. How much value is added by customer data relative to the stand-alone value of the offering?

The higher the value added, the greater the chance that it will create a lasting edge. Let's look at a business where the value of customer data is very high: Mobileye, the leading provider of advanced driver-assistance systems (ADAS), which include collision-prevention and lane-departure warnings for vehicles. Mobileye sells its systems mainly to car manufacturers, which test them extensively before incorporating them into their products. It's crucial for the systems to be fail-safe, and the testing data is essential to improving their accuracy. By

gathering it from dozens of its customers, Mobileye has been able to raise the accuracy of its ADAS to 99.99%.

Conversely, the value of learning from customers is relatively low for makers of smart televisions. Some now include software that can provide personalized recommendations for shows or movies based on an individual's viewing habits as well as what's popular with other users. So far, consumers don't care much about this feature (which is also offered by streaming service providers such as Amazon and Netflix). They largely consider TV size, picture quality, ease of use, and durability when making purchasing decisions. If learning from customers was a bigger factor, perhaps the smart TV business would be less competitive.

2. How quickly does the marginal value of data-enabled learning drop off?

In other words, how soon does the company reach a point where additional customer data no longer enhances the value of an offering? The more slowly the marginal value decreases, the stronger the barrier is. Note that when answering this question, you should judge the value of the learning by customers' willingness to pay and not by some other application-specific measure, such as the percentage of chat-bot queries that could be answered correctly or the fraction of times a movie recommendation was clicked on.

Let's say you graphed the accuracy of Mobileye's ADAS as a function of customer usage (total miles driven by car manufacturers testing it) and found that a few manufacturers and a moderate level of testing would be sufficient to achieve, say, 90% accuracy—but that a lot more testing with a bigger set of car manufacturers would be needed to get to 99%, let alone 99.99%. Interpreting that to mean that the customer data's marginal value was rapidly decreasing would, of course, be incorrect: The value of the additional 9-percentage-point (or even a 0.99-point) improvement in accuracy remains extremely high, given the life-or-death implications. It would be difficult for any individual car manufacturer—even the largest one—to generate the necessary amount of data on its own or for any potential Mobileye competitors to replicate the data. That's why Mobileye was able to carve out a dominant position in the ADAS market, making it a highly attractive acquisition for Intel, which bought it for $15 billion in 2017.

When the marginal value of learning from customer data remains high even after a very large customer base has been acquired, products and services tend to have significant competitive advantages. You can see this with systems designed to predict rare diseases (such as those offered by RDMD) and online search engines such as Baidu and Google. Although Microsoft has invested many years and billions of dollars in Bing, it has been unable to shake Google's dominance in search. Search engines and disease-prediction systems all

need huge amounts of user data to provide consistently reliable results.

A counterexample of a business where the marginal value of user data drops off quickly is smart thermostats. These products need only a few days to learn users' temperature preferences throughout the day. In this context data-enabled learning can't provide much competitive advantage. Although it launched the first smart thermostats that learn from customer behavior in 2011, Nest (acquired by Google in 2014) now faces significant competition from players such as Ecobee and Honeywell.

3. How fast does the relevance of the user data depreciate? If the data becomes obsolete quickly, then all other things being equal, it will be easier for a rival to enter the market, because it doesn't need to match the incumbent's years of learning from data.

All the data Mobileye has accumulated over the years from car manufacturers remains valuable in the current versions of its products. So does the data on search-engine users that Google has collected over decades. Although searches for some terms may become rare over time while searches for new ones might start appearing more frequently, having years of historical search data is of undeniable value in serving today's users. Their data's low depreciation rate helps explain why

both Mobileye and Google Search have proved to be very resilient businesses.

With casual social games for computers and mobile devices, however, the value of learning from user data tends to decrease quickly. In 2009 this market took off when Zynga introduced its highly successful FarmVille game. While the company was famous for relying heavily on user-data analytics to make design decisions, it turned out that the insights learned from one game did not transfer very well to the next: Casual social games are subject to fads, and user preferences shift quickly over time, making it difficult to build sustainable data-driven competitive advantages. After a few more successes, including FarmVille 2 and CityVille, Zynga stopped producing new hits, and in 2013 it lost nearly half its user base. It was superseded by game makers like Supercell (Clash of Clans) and Epic Games (Fortnite). After reaching a peak of $10.4 billion in 2012, Zynga's market value languished below $4 billion for most of the next six years.

4. Is the data proprietary—meaning it can't be purchased from other sources, easily copied, or reverse-engineered?

Having unique customer data with few or no substitutes is critical to creating a defensible barrier. Consider Adaviv, a Boston-area start-up we've invested in, which offers a crop-management system that allows growers (now primarily of cannabis) to continuously monitor individual plants. The

system relies on AI, computer-vision software, and a proprietary data-annotation technique to track plant biometrics not visible to the human eye, such as early signs of disease or lack of adequate nutrients. It then translates the data into insights that growers can use to prevent disease outbreaks and improve yields. The more growers Adaviv serves, the broader the range of variants, agricultural conditions, and other factors it can learn about, and the greater the accuracy of its predictions for new and existing customers. Contrast its situation with that of spam-filter providers, which can acquire user data relatively cheaply. That helps explain the existence of dozens of such providers.

It's important to keep in mind that technological progress can undermine a position based on unique or proprietary data. A case in point is speech-recognition software. Historically, users needed to train the software to understand their individual voices and speech patterns, and the more a person used it, the more accurate it became. This market was dominated by Nuance's Dragon solutions for many years. However, the past decade has seen rapid improvements in speaker-independent speech-recognition systems, which can be trained on publicly available sets of speech data and take minimal or no time to learn to understand a new speaker's voice. These advances have allowed many companies to provide new speech-recognition applications (automated customer service over the phone, automated meeting transcript services, virtual

assistants), and they're putting increasing pressure on Nuance in its core markets.

5. How hard is it to imitate product improvements that are based on customer data?

Even when the data is unique or proprietary and produces valuable insights, it's difficult to build a durable competitive advantage if the resulting enhancements can be copied by competitors without similar data.

A couple of factors affect companies' ability to overcome this challenge. One is whether the improvements are hidden or deeply embedded in a complex production process, making them hard to replicate. Pandora, the music-streaming service, benefits from this barrier. Its offering leveraged the firm's proprietary Music Genome Project, which categorized millions of songs on the basis of some 450 attributes, allowing Pandora to customize radio stations to individual users' preferences. The more a user listens to his or her stations and rates songs up or down, the better Pandora can tailor musical selections to that user. Such customization cannot be easily imitated by any rival because it is deeply tied to the Music Genome Project. In contrast, the design improvements based on learning from the customer use of many office-productivity software products—such as Calendly for coordinating calendars and Doodle for polling people about meeting times—can be easily observed

and copied. That's why dozens of companies offer similar software.

The second factor is how quickly the insights from customer data change. The more rapidly they do so, the harder they are for others to imitate. For example, many design features of the Google Maps interface can be easily copied (and they have been, by Apple Maps, among others). But a key part of Google Maps' value is its ability to predict traffic and recommend optimal routes, which is much harder to copy because it leverages real-time user data that becomes obsolete within minutes. Only companies with similarly large user bases (such as Apple in the United States) can hope to replicate that feature. Apple Maps is closing the gap with Google Maps in the United States, but not in countries where Apple has a relatively small user base.

6. Does the data from one user help improve the product for the same user or for others?

Ideally, it will do both, but the difference between the two is important. When data from one user improves the product for that person, the firm can individually customize it, creating switching costs. When data from one user improves the product for other users, this can—but may not—create network effects. Both kinds of enhancements help provide a barrier to entry, but the former makes *existing* customers very sticky,

whereas the latter provides a key advantage in competing for *new* customers.

For example, Pandora was the first big player in digital music streaming but then fell behind Spotify and Apple Music, which are still growing. As we noted, Pandora's main selling point is that it can tailor stations to each user's tastes. But learning across users is very limited: An individual user's up-or-down votes allow Pandora to identify music attributes that the user likes and then serve that person songs sharing those attributes. In contrast, Spotify focused a lot more on providing users with sharing and discovery features, such as the ability to search and listen to other people's stations, thereby creating direct network effects and luring additional customers. Pandora's service remains available only in the United States (where it has a base of loyal users), while Spotify and Apple Music have become global players. And though Pandora was acquired by Sirius XM for $3.5 billion in February 2019, Spotify became a public company in April 2018 and as of early November 2019 was worth $26 billion. Clearly, customization based on learning from an individual user's data helps keep existing customers locked in, but it doesn't lead to the type of exponential growth that network effects produce.

7. How fast can the insights from user data be incorporated into products?

Rapid learning cycles make it hard for competitors to catch up, especially if multiple product-improvement cycles occur during the average customer's contract. But when it takes years or successive product generations to make enhancements based on the data, competitors have more of a chance to innovate in the interim and start collecting their own user data. So the competitive advantage from customer data is stronger when the learning from *today's* customers translates into more-frequent improvements of the product for those same customers rather than just for *future* customers of the product or service. Several of the product examples we've discussed already—maps, search engines, and AI-based crop-management systems—can be quickly updated to incorporate the learning from current customers.

A counterexample is offered by direct online lenders, such as LendUp and LendingPoint, which learn how to make better loan decisions by examining users' repayment history and how it correlates with various aspects of users' profiles and behavior. Here, the only learning that is relevant to *current* borrowers is that from *previous* borrowers, which is already reflected in the contracts and rates that current borrowers are offered. There's no reason for borrowers to care about any future learning that the lender may benefit from, since their existing contracts won't be affected. For that reason, customers don't worry about how many other borrowers will sign up when deciding whether to take a loan from a particular lender.

Existing borrowers might prefer to stick with their current lenders, which know them better than other lenders do, but the market for new borrowers remains very competitive.

Does Data Confer Network Effects?

The answers to questions 6 and 7 will tell you whether data-enabled learning will create true network effects. When learning from one customer translates into a better experience for other customers *and* when that learning can be incorporated into a product fast enough to benefit its current users, customers will care about how many other people are adopting the product. The mechanism at work here is very similar to the one underlying network effects with online platforms. The difference is that platform users prefer to join bigger networks because they want more people to interact with, not because more users generate more insights that improve products.

Let's look at Google Maps again. Drivers use it in part because they expect many others to employ it too, and the more traffic data the software gathers from them, the better its predictions on road conditions and travel times. Google Search and Adaviv's AI-based crop-management system also enjoy data-enabled network effects.

Like regular network effects, data-enabled ones can create barriers to entry. Both types of effects present a huge cold-start,

or chicken-or-egg, challenge: Businesses aiming to build regular network effects need to attract some minimum number of users to get the effects started, and those aiming to achieve data-enabled network effects need some initial amount of data to start the virtuous cycle of learning.

Despite these similarities, regular network effects and data-enabled network effects have key differences, and they tend to make advantages based on the regular ones stronger. First, the cold-start problem is usually less severe with data-enabled network effects, because buying data is easier than buying customers. Often, alternative sources of data, even if not perfect, can significantly level the playing field by removing the need for a big customer base.

Second, to produce lasting data-enabled network effects, the firm has to work constantly to learn from customer data. In contrast, as Intuit cofounder Scott Cook has often said, "products that benefit from [regular] network effects get better while I sleep." With regular network effects, interactions between customers (and possibly with third-party providers of complementary offerings) create value even if the platform stops innovating. Even if a new social network offered users objectively better features than Facebook does (for instance, better privacy protection), it would still have to contend with Facebook's powerful network effects—users want to be on the same social platform as most other users.

Third, in many cases nearly all the benefits of learning from customer data can be achieved with relatively low numbers of customers. And in some applications (like speech recognition), dramatic improvements in AI will reduce the need for customer data to the point where the value of data-enabled learning might disappear completely. Regular network effects, on the other hand, extend further and are more resilient: An additional customer still typically enhances value for existing customers (who can interact or transact with him or her), even when the number of existing customers is already very large.

As even mundane consumer products become smart and connected—new kinds of clothing, for instance, can now react to weather conditions and track mileage and vital signs—data-enabled learning will be used to enhance and personalize more and more offerings. However, their providers won't build strong competitive positions unless the value added by customer data is high and lasting, the data is proprietary and leads to product improvements that are hard to copy, or the data-enabled learning creates network effects.

In the decades ahead, improving offerings with customer data will be a prerequisite for staying in the game, and it may give incumbents an edge over new entrants. But in most cases it will not generate winner-take-all dynamics. Instead, the most valuable and powerful businesses for the foreseeable

future will be those that are both built on regular network effects and enhanced by data-enabled learning, like Alibaba's and Amazon's marketplaces, Apple's App Store, and Facebook's social networks.

Originally published in January–February 2020. Reprint R2001G

Your Approach to Hiring Is All Wrong

by Peter Cappelli

BUSINESSES HAVE NEVER done as much hiring as they do today. They've never spent as much money doing it. And they've never done a worse job of it.

For most of the post–World War II era, large corporations went about hiring this way: Human resources experts prepared a detailed *job analysis* to determine what tasks the job required and what attributes a good candidate should have. Next they did a *job evaluation* to determine how the job fit into the organizational chart and how much it should pay, especially compared with other jobs. Ads were posted, and applicants applied. Then came the task of sorting through the applicants. That included skills tests, reference checks, maybe personality and IQ tests, and extensive interviews to learn more about them as people. William H. Whyte, in *The Organization Man,* described this process as going on for as long as a week before the winning candidate was offered the job. The vast majority of non-entry-level openings were filled from within.

Today's approach couldn't be more different. Census data shows, for example, that the majority of people who took a new job last year weren't searching for one: Somebody came and got them. Companies seek to fill their recruiting funnel with as many candidates as possible, especially "passive candidates," who aren't looking to move. Often employers advertise jobs that don't exist, hoping to find people who might be useful later on or in a different context.

The recruiting and hiring function has been eviscerated. Many U.S. companies—about 40%, according to research by Korn Ferry —have outsourced much if not all of the hiring process to "recruitment process outsourcers," which in turn use subcontractors, typically in India and the Philippines. The subcontractors scour LinkedIn and social media to find potential candidates. They sometimes contact them directly to see whether they can be persuaded to apply for a position and negotiate the salary they're willing to accept. (The recruiters get incentive pay if they negotiate the amount down.) To hire programmers, for example, these subcontractors can scan websites that programmers might visit, trace their "digital exhaust" from cookies and other user-tracking measures to identify who they are, and then examine their curricula vitae.

At companies that still do their own recruitment and hiring, managers trying to fill open positions are largely left to figure out what the jobs require and what the ads should say. When applications come—always electronically—applicant-tracking

software sifts through them for key words that the hiring managers want to see. Then the process moves into the Wild West, where a new industry of vendors offers an astonishing array of smart-sounding tools that claim to predict who will be a good hire. They use voice recognition, body language, clues on social media, and especially machine learning algorithms—everything but tea leaves. Entire publications are devoted to what these vendors are doing.

The big problem with all these new practices is that we don't know whether they actually produce satisfactory hires. Only about a third of U.S. companies report that they monitor whether their hiring practices lead to good employees; few of them do so carefully, and only a minority even track cost per hire and time to hire. Imagine if the CEO asked how an advertising campaign had gone, and the response was "We have a good idea how long it took to roll out and what it cost, but we haven't looked to see whether we're selling more."

Hiring talent remains the number one concern of CEOs in the most recent Conference Board Annual Survey; it's also the top concern of the entire executive suite. PwC's 2017 CEO survey reports that chief executives view the unavailability of talent and skills as the biggest threat to their business. Employers also spend an enormous amount on hiring—an average of $4,129 per job in the United States, according to Society for Human Resource Management estimates, and many times that amount for managerial roles—and the United States fills a staggering 66

million jobs a year. Most of the $20 billion that companies spend on human resources vendors goes to hiring.

Why do employers spend so much on something so important while knowing so little about whether it works?

Idea in Brief

The Problem

Employers continue to hire at a high rate and spend enormous sums to do it. But they don't know whether their approaches are effective at finding and selecting good candidates.

The Root Causes

Businesses focus on external candidates and don't track the results of their approaches. They often use outside vendors and high-tech tools that are unproven and have inherent flaws.

The Solution

Return to filling most positions by promoting from within. Measure the results produced by vendors and new tools, and be on the lookout for discrimination and privacy violations.

Where the Problem Starts

Survey after survey finds employers complaining about how difficult hiring is. There may be many explanations, such as their having become very picky about candidates, especially in the slack labor market of the Great Recession. But clearly they are hiring much more than at any other time in modern history, for two reasons.

The first is that openings are now filled more often by hiring from the outside than by promoting from within. In the era of lifetime employment, from the end of World War II through the 1970s, corporations filled roughly 90% of their vacancies through promotions and lateral assignments. Today the figure is a third or less. When they hire from outside, organizations don't have to pay to train and develop their employees. Since the restructuring waves of the early 1980s, it has been relatively easy to find experienced talent outside. Only 28% of talent acquisition leaders today report that internal candidates are an important source of people to fill vacancies—presumably because of less internal development and fewer clear career ladders.

Less promotion internally means that hiring efforts are no longer concentrated on entry-level jobs and recent graduates. (If you doubt this, go to the "careers" link on any company website and look for a job opening that doesn't require prior experience.) Now companies must be good at hiring across most levels, because the candidates they want are already doing the job somewhere else. These people don't need training, so they may be ready to contribute right away, but they are much harder to find.

The second reason hiring is so difficult is that retention has become tough: Companies hire from their competitors and vice versa, so they have to keep replacing people who leave. Census and Bureau of Labor Statistics data shows that 95% of hiring is done to fill existing positions. Most of those vacancies are caused

by voluntary turnover. LinkedIn data indicates that the most common reason employees consider a position elsewhere is career advancement—which is surely related to employers' not promoting to fill vacancies.

The root cause of most hiring, therefore, is drastically poor retention. Here are some simple ways to fix that:

Track the percentage of openings filled from within

An adage of business is that we manage what we measure, but companies don't seem to be applying that maxim to tracking hires. Most are shocked to learn how few of their openings are filled from within—is it really the case that their people can't handle different and bigger roles?

Require that all openings be posted internally

Internal job boards were created during the dot-com boom to reduce turnover by making it easier for people to find new jobs within their existing employer. Managers weren't even allowed to know if a subordinate was looking to move within the company, for fear that they would try to block that person and he or she would leave. But during the Great Recession employees weren't quitting, and many companies slid back to the old model whereby managers could prevent their subordinates from moving internally. J.R. Keller, of Cornell University, has found that when managers could fill a vacancy with someone they already had in mind, they ended up with employees who performed more poorly than those hired when the job had been

posted and anyone could apply. The commonsense explanation for this is that few enterprises really know what talent and capabilities they have.

Recognize the costs of outside hiring

In addition to the time and effort of hiring, my colleague Matthew Bidwell found, outside hires take three years to perform as well as internal hires in the same job, while internal hires take seven years to earn as much as outside hires are paid. Outside hiring also causes current employees to spend time and energy positioning themselves for jobs elsewhere. It disrupts the culture and burdens peers who must help new hires figure out how things work.

None of this is to suggest that outside hiring is necessarily a bad idea. But unless your company is a Silicon Valley gazelle, adding new jobs at a furious pace, you should ask yourself some serious questions if most of your openings are being filled from outside.

A different approach for dealing with retention (which seems creepy to some) is to try to determine who is interested in leaving and then intervene. Vendors like Jobvite comb social media and public sites for clues, such as LinkedIn profile updates. Measuring “flight risk” is one of the most common goals of companies that do their own sophisticated HR analytics. This is reminiscent of the early days of job boards, when employers would try to find out who was posting résumés and either punish them or embrace them, depending on leadership’s mood.

Whether companies should be examining social media content in relation to hiring or any other employment action is a challenging ethical question. On one hand, the information is essentially public and may reveal relevant information. On the other hand, it is invasive, and candidates are rarely asked for permission to scrutinize their information. Hiring a private detective to shadow a candidate would also gather public information that might be relevant, yet most people would view it as an unacceptable invasion of privacy.

Protecting Against Discrimination

FINDING OUT WHETHER YOUR PRACTICES result in good hires is not only basic to good management but the only real defense against claims of adverse impact and discrimination. Other than white males under age 40 with no disabilities or work-related health problems, workers have special protections under federal and state laws against hiring practices that may have an adverse impact on them. As a practical matter, that means if members of a particular group are less likely to be recruited or hired, the employer must show that the hiring process is not discriminatory.

The only defense against evidence of adverse impact is for the employer to show that its hiring practices are valid—that is, they predict who will be a good employee in meaningful and statistically significant ways—and that no alternative would predict as well with less adverse impact. That analysis must be conducted with data on the employer's own applicants and hires. The fact that the vendor that sold you the test you use has evidence that it was valid in other contexts is not sufficient.

The Hiring Process

When we turn to hiring itself, we find that employers are missing the forest for the trees: Obsessed with new technologies and driving down costs, they largely ignore the ultimate goal: making the best possible hires. Here's how the process should be revamped:

Don't post "phantom jobs"

It costs nothing to post job openings on a company website, which are then scooped up by Indeed and other online companies and pushed out to potential job seekers around the world. Thus it may be unsurprising that some of these jobs don't really exist. Employers may simply be fishing for candidates. ("Let's see if someone really great is out there, and if so, we'll create a position for him or her.") Often job ads stay up even after positions have been filled, to keep collecting candidates for future vacancies or just because it takes more effort to pull the ad down than to leave it up. Sometimes ads are posted by unscrupulous recruiters looking for résumés to pitch to clients elsewhere. Because these phantom jobs make the labor market look tighter than it really is, they are a problem for economic policy makers as well as for frustrated job seekers. Companies should take ads down when jobs are filled.

Design jobs with realistic requirements

Figuring out what the requirements of a job should be—and the corresponding attributes candidates must have—is a bigger

challenge now, because so many companies have reduced the number of internal recruiters whose function, in part, is to push back on hiring managers' wish lists. ("That job doesn't require 10 years of experience," or "No one with all those qualifications will be willing to accept the salary you're proposing to pay.") My earlier research found that companies piled on job requirements, baked them into the applicant-tracking software that sorted résumés according to binary decisions (yes, it has the key word; no, it doesn't), and then found that virtually no applicants met all the criteria. Trimming recruiters, who have expertise in hiring, and handing the process over to hiring managers is a prime example of being penny-wise and pound-foolish.

Reconsider your focus on passive candidates

The recruiting process begins with a search for experienced people who aren't looking to move. This is based on the notion that something may be wrong with anyone who wants to leave his or her current job. (Of the more than 20,000 talent professionals who responded to a LinkedIn survey in 2015, 86% said their recruiting organizations focused "very much so" or "to some extent" on passive candidates; I suspect that if anything, that number has since grown.) Recruiters know that the vast majority of people are open to moving at the right price: Surveys of employees find that only about 15% are *not* open to moving. As the economist Harold Demsetz said when asked by a competing university if he was happy working where he was: "Make me unhappy."

Fascinating evidence from the LinkedIn survey cited above shows that although self-identified "passive" job seekers are different from "active" job seekers, it's not in the way we might think. The number one factor that would encourage the former to move is more money. For active candidates the top factor is better work and career opportunities. More active than passive job seekers report that they are passionate about their work, engaged in improving their skills, and reasonably satisfied with their current jobs. They seem interested in moving because they are ambitious, not because they want higher pay.

Employers spend a vastly disproportionate amount of their budgets on recruiters who chase passive candidates, but on average they fill only 11% of their positions with individually targeted people, according to research by Gerry Crispin and Chris Hoyt, of CareerXroads. I know of no evidence that passive candidates become better employees, let alone that the process is cost-effective. If you focus on passive candidates, think carefully about what that actually gets you. Better yet, check your data to find out.

Understand the limits of referrals

The most popular channel for finding new hires is through employee referrals; up to 48% come from them, according to LinkedIn research. It seems like a cheap way to go, but does it produce better hires? Many employers think so. It's hard to know whether that's true, however, given that they don't check. And research by Emilio Castilla and colleagues suggests otherwise:

They find that when referrals work out better than other hires, it's because their referrers look after them and essentially onboard them. If a referrer leaves before the new hire begins, the latter's performance is no better than that of nonreferrals, which is why it makes sense to pay referral bonuses six months or so after the person is hired—if he or she is still there.

A downside to referrals, of course, is that they can lead to a homogeneous workforce, because the people we know tend to be like us. This matters greatly for organizations interested in diversity, since recruiting is the only avenue allowed under U.S. law to increase diversity in a workforce. The Supreme Court has ruled that demographic criteria cannot be used even to break ties among candidates.

Measure the results

Few employers know which channel produces the best candidates at the lowest cost because they don't track the outcomes. Tata is an exception: It has long done what I advocate. For college recruiting, for example, it calculates which schools send it employees who perform the best, stay the longest, and are paid the lowest starting wage. Other employers should follow suit and monitor recruiting channels and employees' performance to identify which sources produce the best results.

Persuade fewer people to apply

The hiring industry pays a great deal of attention to “the funnel,” whereby readers of a company’s job postings become applicants, are interviewed, and ultimately are offered jobs. Contrary to the popular belief that the U.S. job market is extremely tight right now, most jobs still get lots of applicants. Recruiting and hiring consultants and vendors estimate that about 2% of applicants receive offers. Unfortunately, the main effort to improve hiring—virtually always aimed at making it faster and cheaper—has been to shovel more applicants into the funnel. Employers do that primarily through marketing, trying to get out the word that they are great places to work. Whether doing this is a misguided way of trying to attract better hires or just meant to make the organization feel more desirable isn’t clear.

Much better to go in the other direction: Create a smaller but better-qualified applicant pool to improve the yield. Here’s why: Every applicant costs you money—especially now, in a labor market where applicants have started to “ghost” employers, abandoning their applications midway through the process. Every application also exposes a company to legal risk, because the company has obligations to candidates (not to discriminate, for example) just as it does to employees. And collecting lots of applicants in a wide funnel means that a great many of them won’t fit the job or the company, so employers have to rely on the next step of the hiring process—selection—to weed them out. As we will see, employers aren’t good at that.

The grass is always greener …

Organizations are much more interested in external talent than in their own employees to fill vacancies.

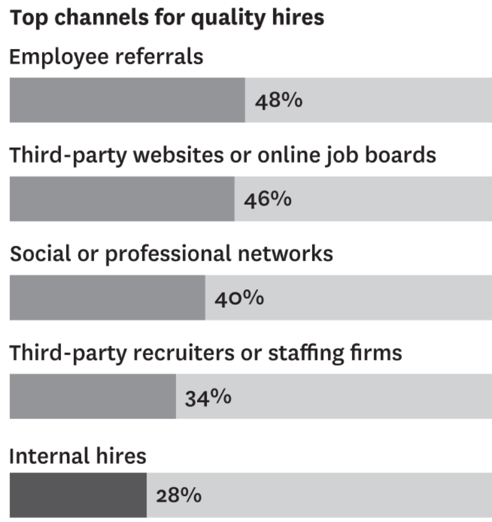

Source: LinkedIn, based on a 2017 survey of 3,973 talent-acquisition decision makers who work in corporate HR departments and are LinkedIn members.

Once people are candidates, they may not be completely honest about their skills or interests—because they want to be hired—and employers' ability to find out the truth is limited. More than a generation ago the psychologist John Wanous proposed giving applicants a realistic preview of what the job is like. That still makes sense as a way to head off those who would end up being unhappy in the job. It's not surprising that Google has found a way to do this with gamification: Job seekers see what the work would be like by playing a game version of it. Marriott has done the same, even for low-level employees. Its My Marriott Hotel game targets young people in developing countries who may have had little experience in hotels to show

them what it's like and to steer them to the recruiting site if they score well on the game. The key for any company, though, is that the preview should make clear what is difficult and challenging about the work as well as why it's fun so that candidates who don't fit won't apply.

It should be easy for candidates to learn about a company and a job, but making it really easy to apply, just to fill up that funnel, doesn't make much sense. During the dot-com boom Texas Instruments cleverly introduced a preemployment test that allowed applicants to see their scores before they applied. If their scores weren't high enough for the company to take their applications seriously, they tended not to proceed, and the company saved the cost of having to process their applications.

If the goal is to get better hires in a cost-effective manner, it's more important to scare away candidates who don't fit than to jam more candidates into the recruiting funnel.

Test candidates' standard skills

How to determine which candidates to hire—what predicts who will be a good employee—has been rigorously studied at least since World War I. The personnel psychologists who investigated this have learned much about predicting good hires that contemporary organizations have since forgotten, such as that neither college grades nor unstructured sequential interviews (hopping from office to office) are a good predictor, whereas past performance is.

Since it can be difficult (if not impossible) to glean sufficient information about an outside applicant's past performance, what other predictors are good? There is remarkably little consensus even among experts. That's mainly because a typical job can have so many tasks and aspects, and different factors predict success at different tasks.

There is general agreement, however, that testing to see whether individuals have standard skills is about the best we can do. Can the candidate speak French? Can she do simple programming tasks? And so forth. But just doing the tests is not enough. The economists Mitchell Hoffman, Lisa B. Kahn, and Danielle Li found that even when companies conduct such tests, hiring managers often ignore them—and when they do, they get worse hires. The psychologist Nathan Kuncel and colleagues discovered that even when hiring managers use objective criteria and tests, applying their own weights and judgment to those criteria leads them to pick worse candidates than if they had used a standard formula. Only 40% of employers, however, do any tests of skills or general abilities, including IQ. What are they doing instead? Seventy-four percent do drug tests, including for marijuana use; even employers in states where recreational use is now legal still seem to do so.

Be wary of vendors bearing high-tech gifts

Into the testing void has come a new group of entrepreneurs who either are data scientists or have them in tow. They bring a fresh approach to the hiring process—but often with little

understanding of how hiring actually works. John Sumser, of HRExaminer, an online newsletter that focuses on HR technology, estimates that on average, companies get five to seven pitches *every day*—almost all of them about hiring—from vendors using data science to address HR issues. These vendors have all sorts of cool-sounding assessments, such as computer games that can be scored to predict who will be a good hire. We don't know whether any of these actually lead to better hires, because few of them are validated against actual job performance. That aside, these assessments have spawned a counterwave of vendors who help candidates learn how to score well on them. Lloyds Bank, for example, developed a virtual-reality-based assessment of candidate potential, and JobTestPrep offers to teach potential candidates how to do well on it. Especially for IT and technical jobs, cheating on skills tests and even video interviews (where colleagues off camera give help) is such a concern that eTeki and other specialized vendors help employers figure out who is cheating in real time.

Revamp your interviewing process

The amount of time employers spend on interviews has almost doubled since 2009, according to research from Glassdoor. How much of that increase represents delays in setting up those interviews is impossible to tell, but it provides at least a partial explanation for why it takes longer to fill jobs now. Interviews are arguably the most difficult technique to get right, because interviewers should stick to questions that predict good hires—

mainly about past behavior or performance that's relevant to the tasks of the job—and ask them consistently across candidates. Just winging it and asking whatever comes to mind is next to useless.

More important, interviews are where biases most easily show up, because interviewers do usually decide on the fly what to ask of whom and how to interpret the answer. Everyone knows some executive who is absolutely certain he knows the one question that will really predict good candidates ("If you were stranded on a desert island …"). The sociologist Lauren Rivera's examination of interviews for elite positions, such as those in professional services firms, indicates that hobbies, particularly those associated with the rich, feature prominently as a selection criterion.

Interviews are most important for assessing "fit with our culture," which is the number one hiring criterion employers report using, according to research from the Rockefeller Foundation. It's also one of the squishiest attributes to measure, because few organizations have an accurate and consistent view of their own culture—and even if they do, understanding what attributes represent a good fit is not straightforward. For example, does the fact that an applicant belonged to a fraternity reflect experience working with others or elitism or bad attitudes toward women? Should it be completely irrelevant? Letting someone with no experience or training make such calls is a recipe for bad hires and, of course, discriminatory behavior.

Think hard about whether your interviewing protocols make any sense and resist the urge to bring even more managers into the interview process.

Recognize the strengths and weaknesses of machine learning models

Culture fit is another area into which new vendors are swarming. Typically they collect data from current employees, create a machine learning model to predict the attributes of the best ones, and then use that model to hire candidates with the same attributes.

As with many other things in this new industry, that sounds good until you think about it; then it becomes replete with problems. Given the best performers of the past, the algorithm will almost certainly include *white* and *male* as key variables. If it's restricted from using that category, it will come up with attributes associated with being a white male, such as playing rugby.

Machine learning models do have the potential to find important but previously unconsidered relationships. Psychologists, who have dominated research on hiring, have been keen to study attributes relevant to their interests, such as personality, rather than asking the broader question "What identifies a potential good hire?" Their results gloss over the fact that they often have only a trivial ability to predict who will be a good performer, particularly when many factors are involved. Machine learning, in contrast, can come up with highly

predictive factors. Research by Evolv, a workforce analytics pioneer (now part of Cornerstone OnDemand), found that expected commuting distance for the candidate predicted turnover very well. But that's not a question the psychological models thought to ask. (And even that question has problems.)

The advice on selection is straightforward: Test for skills. Ask assessments vendors to show evidence that they can actually predict who the good employees will be. Do fewer, more-consistent interviews.

Expanding the Pool: How Goldman Sachs Changed the Way It Recruits

by Dane E. Holmes

Goldman Sachs is a people-centric business—every day our employees engage with our clients to find solutions to their challenges. As a consequence, hiring extraordinary talent is vital to our success and can never be taken for granted. In the wake of the 2008 financial crisis we faced a challenge that was, frankly, relatively new to our now 150-year-old firm. For decades investment banking had been one of the most sought-after, exciting, and fast-growing industries in the world. That made

sense—we were growing by double digits and had high returns, which meant that opportunity and reward were in great supply. However, the crash took some of the sheen off our industry; both growth and returns moderated. And simultaneously, the battle for talent intensified—within and outside our industry. Many of the candidates we were pursuing were heading off to Silicon Valley, private equity, or start-ups. Furthermore, we were no longer principally looking for a specialized cadre of accounting, finance, and economics majors: New skills, especially coding, were in huge demand at Goldman Sachs—and pretty much everywhere else. The wind had shifted from our backs to our faces, and we needed to respond.

Not long ago the firm relied on a narrower set of factors for identifying the "best" students, such as school, GPA, major, leadership roles, and relevant experience—the classic résumé topics. No longer. We decided to replace our hiring playbook with emerging best practices for assessment and recruitment, so we put together a task force of senior business leaders, PhDs in industrial and organizational psychology, data scientists, and experts in recruiting. Some people asked, "Why overhaul a recruiting process that has proved so successful?" and "Don't you already have many more qualified applicants than available jobs?" These were reasonable questions. But often staying successful is about learning and changing rather than sticking to the tried-and-true.

Each year we hire up to 3,000 summer interns and nearly as many new analysts directly from campuses. In our eyes, these are the firm's future leaders, so it made sense to focus our initial reforms there. They involved two major additions to our campus recruiting strategy—video interviews and structured interviewing.

Asynchronous Video Interviews

Traditionally we had flown recruiters and business professionals to universities for first-round interviews. The schools would give us a set date and number of time slots to meet with students. That is most definitely not a scalable model. It restricted us to a smaller number of campuses and only as many students as we could squeeze into a limited schedule. It also meant that we tended to focus on top-ranked schools. How many qualified candidates were at a school became more important than who were the most talented students regardless of their school. However, we knew that candidates didn't have to attend Harvard, Princeton, or Oxford to excel at Goldman Sachs—our leadership ranks were already rich with people from other schools. What's more, as we've built offices in new cities and geographic locations, we've needed to recruit at more schools located in those areas. Video interviews allow us to do that.

At a time when companies were just beginning to experiment with digital interviewing, we decided to use "asynchronous" video interviews—in which candidates record their answers to

interview questions—for all first-round interactions with candidates. Our recruiters record standardized questions and send them to students, who have three days to return videos of their answers. This can be done on a computer or a mobile device. Our recruiters and business professionals review the videos to narrow the pool and then invite the selected applicants to a Goldman Sachs office for final-round, in-person interviews. (To create the video platform, we partnered with a company and built our own digital solution around its product.)

This approach has had a meaningful impact in two ways. First, with limited effort, we can now spend more time getting to know the people who apply for jobs at Goldman Sachs. In 2015, the year before we rolled out this platform, we interviewed fewer than 20% of all our campus applicants; in 2018 almost 40% of the students who applied to the firm participated in a first-round interview. Second, we now encounter talent from places we previously didn't get to. In 2015 we interviewed students from 798 schools around the world, compared with 1,268 for our most recent incoming class. In the United States, where the majority of our student hires historically came from "target schools," the opposite is now true. The top of our recruiting funnel is wider, and the output is more diverse.

Being a people-driven business, we have worked hard to ensure that the video interviews don't feel cold and impersonal. They are only one component of a broader process that makes up the Goldman Sachs recruitment experience. We still regularly

send Goldman professionals to campuses to engage directly with students at informational sessions, "coffee chats," and other recruiting events. But now our goal is much more to share information than to assess candidates, because we want people to understand the firm and what it offers before they tell us why they want an internship or a job.

We also want them to be as well prepared as possible for our interview process. Our goal is a level playing field. To help achieve it, we've created tip sheets and instructions on preparing for a video interview. Because the platform doesn't allow videos to be edited once they've been recorded, we offer a practice question before the interview begins and a countdown before the questions are asked. We also give students a formal channel for escalating issues should technical problems arise, though that rarely occurs.

We're confident that this approach has created a better experience for recruits. It uses a medium they've grown up with (video), and most important, they can do their interviews when they feel fresh and at a time that works with their schedule. (Our data shows that they prefer Thursday or Sunday night—whereas our previous practice was to interview during working hours.) We suspected that if the process was a turnoff for applicants, we would see a dip in the percentage who accepted our interviews and our offers. That hasn't happened.

Structured Questioning and Assessments

How can you create an assessment process that not only helps select top talent but focuses on specific characteristics associated with success? Define it, structure it, and don't deviate from it. Research shows that structured interviews are effective at assessing candidates and helping predict job performance. So we ask candidates about specific experiences they've had that are similar to situations they may face at Goldman Sachs ("Tell me about a time when you were working on a project with someone who was not completing his or her tasks") and pose hypothetical scenarios they might encounter in the future ("In an elevator, you overhear confidential information about a coworker who is also a friend. The friend approaches you and asks if you've heard anything negative about him recently. What do you do?").

Essentially, we are focused less on past achievements and more on understanding whether a candidate has qualities that will positively affect our firm and our culture. Our structured interview questions are designed to assess candidates on 10 core competencies, including analytical thinking and integrity, which we know correlate with long-term success at the firm. They are evaluated on six competencies in the first round; if they progress, they're assessed on the remaining four during in-person interviews.

We have a rotating library of questions for each competency, along with a rubric for interviewers that explains how to rate responses on a five-point scale from "outstanding" to "poor." We also train our interviewers to conduct structured interviews,

provide them with prep materials immediately before they interview a candidate, and run detailed calibration meetings using all the candidate data we've gathered throughout the recruiting process to ensure that certain interviewers aren't introducing grade inflation (or deflation). We're experimenting with prehire assessment tests to be paired with these interviews; we already offer a technical coding and math exam for applicants to our engineering organization.

We decided not to pilot these changes and instead rolled them out en masse, because we realized that buy-in would come from being able to show results quickly—and because we know that no process is perfect. Indeed, what I love most about our new approach is that we've turned our recruiting department into a laboratory for continuous learning and refinement. With more than 50,000 candidate video recordings, we're now sitting on a treasure trove of data that will help us conduct insightful analyses and answer questions necessary to run our business: Are we measuring the right competencies? Should some be weighted more heavily than others? What about the candidates' backgrounds? Which interviewers are most effective? Does a top-ranked student at a state school create more value for us than an average student from the Ivy League? We already have indications that students recruited from the new schools in our pool perform just as well as students from our traditional ones—and in some cases are more likely to stay longer at the firm.

What's next for our recruiting efforts? We receive almost 500,000 applications each year. From this pool we hire approximately 3%. We believe that many of the other 97% could be very successful at Goldman Sachs. As a result, picking the right 3% is less about just the individual and increasingly about matching the right person to the right role. That match may be made straight out of college or years later. We're experimenting with résumé-reading algorithms that will help candidates identify the business departments best suited to their skills and interests. We're looking at how virtual reality might help us better educate students about working in our offices and in our industry. And we're evaluating various tools and tests to bring even more data into the hiring decision process. Can I imagine a future in which companies rely exclusively on machines and algorithms to rate résumés and interviews? Maybe, for some. But I don't see us ever eliminating the human element at Goldman Sachs; it's too deeply embedded in our culture, in the work we do, and in what we believe drives success.

I'm excited to see where this journey takes us. Our 2019 campus class is shaping up to be the most diverse ever—and it's composed entirely of people who were selected through rigorous, objective assessments. There's no way we aren't better off as a result.

Data Science Can't Fix Hiring (Yet)

by Peter Cappelli

Recruiting managers desperately need new tools, because the existing ones—unstructured interviews, personality tests, personal referrals—aren't very effective. The newest development in hiring, which is both promising and worrying, is the rise of data science–driven algorithms to find and assess job candidates. By my count, more than 100 vendors are creating and selling these tools to companies. Unfortunately, data science—which is still in its infancy when it comes to recruiting and hiring—is not yet the panacea employers hope for.

Vendors of these new tools promise they will help reduce the role that social bias plays in hiring. And the algorithms can indeed help identify good job candidates who would previously have been screened out for lack of a certain education or social pedigree. But these tools may also identify and promote the use of predictive variables that are (or should be) troubling.

Because most data scientists seem to know so little about the context of employment, their tools are often worse than nothing. For instance, an astonishing percentage build their models by simply looking at attributes of the "best performers" in workplaces and then identifying which job candidates have the same attributes. They use anything that's easy to measure: facial expressions, word choice, comments on social media, and so

forth. But a failure to check for any real difference between high-performing and low-performing employees on these attributes limits their usefulness. Furthermore, scooping up data from social media or the websites people have visited also raises important questions about privacy. True, the information can be accessed legally; but the individuals who created the postings didn't intend or authorize them to be used for such purposes. Furthermore, is it fair that something you posted as an undergrad can end up driving your hiring algorithm a generation later?

Another problem with machine learning approaches is that few employers collect the large volumes of data—number of hires, performance appraisals, and so on—that the algorithms require to make accurate predictions. Although vendors can theoretically overcome that hurdle by aggregating data from many employers, they don't really know whether individual company contexts are so distinct that predictions based on data from the many are inaccurate for the one.

Yet another issue is that all analytic approaches to picking candidates are backward looking, in the sense that they are based on outcomes that have already happened. (Algorithms are especially reliant on past experiences in part because building them requires lots and lots of observations—many years' worth of job performance data even for a large employer.) As Amazon learned, the past may be very different from the future you seek. It discovered that the hiring algorithm it had been working on

since 2014 gave lower scores to women—even to attributes associated with women, such as participating in women's studies programs—because historically the best performers in the company had disproportionately been men. So the algorithm looked for people just like them. Unable to fix that problem, the company stopped using the algorithm in 2017. Nonetheless, many other companies are pressing ahead.

The underlying challenge for data scientists is that hiring is simply not like trying to predict, say, when a ball bearing will fail—a question for which any predictive measure might do. Hiring is so consequential that it is governed not just by legal frameworks but by fundamental notions of fairness. The fact that some criterion is associated with good job performance is necessary but not sufficient for using it in hiring.

Take a variable that data scientists have found to have predictive value: commuting distance to the job. According to the data, people with longer commutes suffer higher rates of attrition. However, commuting distance is governed by where you live—which is governed by housing prices, relates to income, and also relates to race. Picking whom to hire on the basis of where they live most likely has an adverse impact on protected groups such as racial minorities.

Unless no other criterion predicts at least as well as the one being used—and that is extremely difficult to determine in machine learning algorithms—companies violate the law if they use hiring criteria that have adverse impacts. Even then, to stay

on the right side of the law, they must show why the criterion creates good performance. That might be possible in the case of commuting time, but—at least for the moment—it is not for facial expressions, social media postings, or other measures whose significance companies cannot demonstrate.

In the end, the drawback to using algorithms is that we're trying to use them on the cheap: building them by looking only at best performers rather than all performers, using only measures that are easy to gather, and relying on vendors' claims that the algorithms work elsewhere rather than observing the results with our own employees. Not only is there no free lunch here, but you might be better off skipping the cheap meal altogether.

The Way Forward

It's impossible to get better at hiring if you can't tell whether the candidates you select become good employees. If you don't know where you're going, any road will take you there. You must have a way to measure which employees are the best ones.

Why is that not getting through to companies? Surveyed employers say the main reason they don't examine whether their practices lead to better hires is that measuring employee performance is difficult. Surely this is a prime example of making the perfect the enemy of the good. Some aspects of performance are not difficult to measure: Do employees quit? Are they absent? Virtually all employers conduct performance appraisals. If you

don't trust them, try something simpler. Ask supervisors, "Do you regret hiring this individual? Would you hire him again?"

Organizations that don't check to see how well their practices predict the quality of their hires are lacking in one of the most consequential aspects of modern business.

Originally published in May–June 2019. Reprint R1903B

How Dual-Career Couples Make It Work

by Jennifer Petriglieri

CAMILLE AND PIERRE MET in their early forties after each one's marriage had ended. Both were deeply committed to their careers and to their new relationship. Camille, an accountant, had felt pressured by her ex-husband to slow her progress toward partnership at her firm. Pierre, a production manager at an automotive company, was embroiled in a bitter divorce from his wife, who had given up her career to accommodate the geographic moves that his required. (As with the other couples I've profiled in this article, these aren't their real names.) Bruised by their past experiences, they agreed to place their careers on an equal footing. Initially things went smoothly, but two years in, Camille began to feel trapped on a professional path that she realized she had chosen because "that was what the smart kids did."

Mindful of their pact, Pierre calmly listened to her doubts and encouraged her to explore alternatives. But as the months

wore on, he began to feel weighed down as he juggled providing emotional support to Camille, navigating their complex family logistics (both had children from their former marriages), and succeeding in his demanding job. When he began to question his own career direction, he wondered how the two of them could manage to change course. They couldn't afford to take time out from work, nor could they take much time to reflect and keep their family and relationship afloat. Frustrated and exhausted, both wondered how they could continue to find meaning and fulfillment in their lives.

Dual-earner couples are on the rise. According to Pew Research, in 63% of couples with children in the United States, for example, both partners work (this figure is slightly higher in the EU). Many of these are *dual-career couples:* Both partners are highly educated, work full-time in demanding professional or managerial jobs, and see themselves on an upward path in their roles. For these couples, as for Pierre and Camille, work is a primary source of identity and a primary channel for ambition. Evidence is mounting from sociological research that when both partners dedicate themselves to work and to home life, they reap benefits such as increased economic freedom, a more satisfying relationship, and a lower-than-average chance of divorce.

Because their working lives and personal lives are deeply intertwined, however, dual-career couples face unique challenges. How do they decide whose job to relocate for, when

it's OK for one partner to make a risky career change, or who will leave work early to pick up a sick child from school? How can they give family commitments—and each other—their full attention while both of them are working in demanding roles? And when one of them wants to undertake a professional reinvention, what does that mean for the other? They must work out these questions together, in a way that lets both thrive in love and work. If they don't, regrets and imbalances quickly build up, threatening to hinder their careers, dissolve their relationship, or both.

Many of these challenges are well recognized, and I've previously written in HBR about how companies can adapt their talent strategies to account for some of them ("Talent Management and the Dual-Career Couple," May–June 2018). But for the couples themselves, little guidance is available. Most advice treats major career decisions as if one is flying solo, without a partner, children, or aging parents to consider. When it's for couples, it focuses on their relationship, not how that intersects with their professional dreams, or it addresses how to balance particular trade-offs, such as careers versus family, or how to prioritize partners' work travel. What couples need is a more comprehensive approach for managing the moments when commitments and aspirations clash.

Idea in Brief

The Problem

When both members of a couple have demanding careers, their work and personal lives are deeply intertwined—and often at odds.

The Transitions

Dual-career couples tend to go through three phases of being particularly vulnerable: when they first learn to work together as a couple; when they experience a midlife reinvention; and in the final stages of their working lives.

The Solution

Couples who communicate at each transition about values, boundaries, and fears have a good chance of being fulfilled both in their relationships and in their careers.

My personal experience in a dual-career couple, and my realization that little systematic academic research had been done in this area, prompted a six-year investigation into the lives of more than 100 dual-career couples, resulting in my book, *Couples That Work*. The people I studied come from around the world, range in age from mid-twenties to mid-sixties, and represent a range of professions, from corporate executive to entrepreneur to worker in the nonprofit sector. (See the sidebar "About the Research.") My research revealed that dual-career couples overcome their challenges by directly addressing deeper psychological and social forces—such as struggles for power and control; personal hopes, fears, and losses; and assumptions and cultural expectations about the roles partners should play in each other's lives and what it means to have a good relationship or career.

About the Research

I STUDIED 113 DUAL-CAREER COUPLES. They ranged in age from 26 to 63, with an even distribution among age groups. The majority of couples—76—were in their first significant partnership. Participants in the study came from 32 countries on four continents, and their ethnic and religious backgrounds reflected this diversity. At the time of the study, roughly 35% resided in North America, 40% in Europe, and 25% in the rest of the world. In 68 of the couples at least one partner had children. Eleven of the couples identified as gay, and the rest as straight. Just under 60% of the participants were pursuing careers in the corporate world. The others were spread roughly equally among the professions (such as medicine, law, and academia), entrepreneurship, government, and the nonprofit sector.

I interviewed the members of each couple separately, asking them about the development of their relationships, their career paths, their interactions as a couple, and their family and friend networks.

I also discovered that three transition points typically occur during dual-career couples' working and love lives, when those forces are particularly strong. It is during these transitions, I found, that some couples craft a way to thrive in love and work, while others are plagued by conflict and regret. By understanding each transition and knowing what questions to ask each other and what traps to avoid, dual-career couples can emerge stronger, fulfilled in their relationships and in their careers.

Transition 1: Working as a Couple

When Jamal and Emily met, in their late twenties, trade-offs were the last thing on their minds. They were full of energy, optimistic, and determined to live life to the fullest. Jamal, a project manager in a civil engineering firm, traveled extensively for work and was given increasingly complex projects to lead, while Emily, who worked at a clothing company, had just been promoted to her first management role. They saw each other mostly on weekends, which they often spent on wilderness hiking adventures. They married 18 months after their first date.

Then, in the space of three months, their world changed dramatically. While Emily was pregnant with their first child, Jamal's boss asked him to run a critical infrastructure project in Mexico. Jamal agreed to spend three weeks out of every month in Mexico City; designating some of his pay raise to extra child care would allow Emily to keep working in Houston, where they lived. But when their daughter, Aisha, was born two weeks early, Jamal was stuck in the Mexico City airport waiting for a flight home. Soon Emily, who was single-handedly managing Aisha, her job, and their home, discovered that the additional child care wasn't enough; she felt overburdened and unappreciated. Jamal was exhausted by the relentless travel and the stress of the giant new project; he felt isolated, incompetent, and guilty.

After many arguments, they settled on what they hoped was a practical solution: Because Jamal earned more, Emily took a smaller project role that she could manage remotely, and she and Aisha joined him in Mexico. But Emily felt disconnected from her company's head office and was passed over for a promotion, and eventually she grew resentful of the arrangement. By the time Jamal's boss began talking about his next assignment, their fighting had become intense.

The first transition that dual-career couples must navigate often comes as a response to the first major life event they face together—typically a big career opportunity, the arrival of a child, or the merger of families from previous relationships. To adapt, the partners must negotiate how to prioritize their careers and divide family commitments. Doing so in a way that lets them both thrive requires an underlying shift: They must move from having parallel, independent careers and lives to having interdependent ones.

My research shows two common traps for couples negotiating their way through their first transition:

Concentrating exclusively on the practical

In the first transition in particular, couples often look for logistical solutions to their challenges, as Jamal and Emily did when they arranged for extra child care and negotiated how many weekends Jamal would be home. This focus is understandable—such problems are tangible, and the

underlying psychological and social tensions are murky and anxiety provoking—but it prolongs the struggle, because those tensions remain unresolved.

Instead of simply negotiating over calendars and to-do lists, couples must understand, share, and discuss the emotions, values, and fears underlying their decisions. Talking about feelings as well as practicalities can help them mitigate and manage them.

Basing decisions primarily on money

Many couples focus on economic gain as they decide where to live, whose career to prioritize, and who will do the majority of the child care. But as sensible (and sometimes unavoidable) as this is, it often means that their decisions end up at odds with their other values and desires.

Few people live for financial gain alone. In their careers they are also motivated by continual learning and being given greater responsibilities. Outside work, they want to spend time with their children and pursue personal interests. Couples may be attracted to a location because of proximity to extended family, the quality of life it affords, or their ability to build a strong community. Basing the decision to move to Mexico on Jamal's higher salary meant that he and Emily ignored their other interests, feeding their discontent.

Couples who are successful discuss the foundations and the structure of their joint path forward. First, they must come to

some agreement on core aspects of their relationship: their values, boundaries, and fears. (See the sidebar "A Guide to Couple Contracting.") Negotiating and finding common ground in these areas helps them navigate difficult decisions because they can agree on criteria in advance. Doing this together is important; couples that make this arrangement work, I found, make choices openly and jointly, rather than implicitly and for each other. The ones I studied who had never addressed their core criteria struggled in later transitions, because those criteria never go away.

A Guide to Couple Contracting

DRAWING ON MY RESEARCH, I've developed a systematic tool to help dual-career couples who are facing any of the three transitions described in this article. I call it couple contracting, because to shape their joint path, partners must address three areas—values, boundaries, and fears—and find common ground in each. Values define the direction of your path, boundaries set its borders, and fears reveal the potential cliffs to avoid on either side. Sharing a clear view in these three domains will make it easier to negotiate and overcome the challenges you encounter together.

First, take some time on your own to write down your thoughts about each of the three areas. Then share your reflections with each other. Listen to and acknowledge each other's responses, resisting any temptation to diminish or discount your partner's fears. Next, note where you have common ground and where your values and boundaries diverge. No couple has perfect overlap in those two areas, but if they are too divergent, negotiate a middle ground. If, for example, one of you could tolerate living apart for a period but the

other could not, you'll need to shape a boundary that works for both of you.

Values

When our choices and actions align with our values, we feel content; when they don't, we feel stressed and unhappy. Openly discussing your values will make it easier to make choices that align with them. For example, if you and your partner know you both greatly value family time, you'll be clear that neither of you should take a job requiring 70-hour workweeks.

Questions to ask each other

What makes you happy and proud? What gives you satisfaction? What makes for a good life?

Boundaries

Setting clear boundaries together allows you to make big decisions more easily. Consider three types of boundaries: place, time, and presence.

Questions to ask each other

Are there places where you'd love to work and live at some point in your life? Are there places you'd prefer to avoid? Understanding that we may sometimes have to put in more hours than we'd like, how much work is too much? How would you feel about our taking jobs in different cities and living apart for a period? For how long? How much work travel is too much, and how will we juggle travel between us?

Fears

Monitoring each other's fears can help you spot trouble and take preventive action before your relationship enters dangerous territory. Many fears are endemic to relationships and careers: You may worry that your partner's family will encroach on your relationship, that over time the two of you will grow apart, that your partner will have an

affair, that you will have to sacrifice your career for your partner's, or that you may not be able to have children. But sharing these fears allows you to build greater empathy and support. If you know that your partner is worried about the role of your parents in your lives, for example, you are more likely to manage the boundary between them and your partnership sensitively. Likewise, if you are interested in a risky career transition but worried that financial commitments would prevent it, you might agree to cut back on family spending in order to build a buffer.

Questions to ask each other

What are your concerns for the future? What's your biggest fear about how our relationship and careers interact? What do you dread might happen in our lives?

Next, couples must discuss how to prioritize their careers and divide family commitments. Striving for 50/50 is not always the best option; neither must one decide to always give the other's career priority.

There are three basic models to consider: (1) In *primary-secondary,* one partner's career takes priority over the other's for the duration of their working lives. The primary person dedicates more time to work and less to the family, and his or her professional commitments (and geographic requirements) usually come before the secondary person's. (2) In *turn taking,* the partners agree to periodically swap the primary and secondary positions. (3) In *double-primary,* they continually juggle two primary careers.

My research shows that couples can feel fulfilled in their careers and relationships whichever model they pursue, as long as it aligns with their values and they openly discuss and explicitly agree on their options. Couples who pursue the third option are often the most successful, although it's arguably the most difficult, precisely because they are forced to address conflicts most frequently.

To work past their deadlock, Emily and Jamal finally discussed what really mattered to them beyond financial success. They identified pursuit of their chosen careers, proximity to nature, and a stable home for Aisha where they could both actively parent her. They admitted their fears of growing apart, and in response agreed to an important restriction: They would live in the same city and would limit work travel to 25% of their time. They agreed to place their geographic boundaries around North America, and Jamal suggested that they both draw circles on a map around the cities where they felt they could make a home and have two careers. Their conversations and mapping exercise eventually brought them to a resolution—and a new start in Atlanta, where they would pursue a double-primary model. Three years later they are progressing in their careers, happy in their family life, and expecting a second child.

Transition 2: Reinventing Themselves

Psychological theory holds that early in life many people follow career and personal paths that conform to the expectations of their parents, friends, peers, and society, whereas in their middle years many feel a pressing need for *individuation,* or breaking free of those expectations to become authors of their own lives. This tends to happen in people's forties, regardless of their relationship status, and is part of a process colloquially known as the midlife crisis.

We tend to think of a midlife crisis mostly in personal terms (a husband leaves his wife, for example, and buys a sports car), but in dual-career couples, the intense focus on professional success means that the partners' job tracks come under scrutiny as well. This combined personal and professional crisis forms the basis of the second transition. Camille and Pierre, whose story began this article, were in the midst of it.

As each partner wrestles with self-redefinition, the two often bump up against long-settled arrangements they have made and the identities, relationship, and careers they have crafted together. Some of those arrangements—whose career takes precedence, for example—may need to be reconsidered to allow one partner to quit a job and explore alternatives. It may be painful to question the choices they made together during the previous transition and have since built their lives around. This can be threatening to a relationship; it's not uncommon for one partner to interpret the other's desire to rethink past career choices as an inclination to rethink the relationship as

well, or even to potentially end it. Couples who handle this transition well find ways to connect with and support each other through what can feel like a very solitary process.

The second transition often begins—as it did for Camille and Pierre—when one partner reexamines a career or life path. That person must reflect on questions such as: What led me to this impasse? Why did I make the choices I made? Who am I? What do I desire from life? Whom do I want to become? He or she should also take time to explore alternative paths, through networking events, job shadowing, secondments, volunteer work, and so forth. Such individual reflection and exploration can lead couples to the first trap of the second transition:

Mistrust and defensiveness

Living with a partner who is absorbed in exploring new paths can feel threatening. Painful questions surface: Why is my partner not satisfied? Is this a career problem or a relationship problem? Am I to blame? Why does he or she need new people? Am I no longer enough? These doubts can lead to mistrust and defensiveness, which may push the exploring partner to withdraw further from the relationship, making the other even more mistrustful and defensive, until eventually the relationship itself becomes an obstacle to individuation, rather than a space for it.

In such a situation, people should first be open about their concerns and let their partners reassure them that the angst is

not about them or the relationship. Next, they should adopt what literary critics call *suspension of disbelief*—that is, faith that the things they have doubts about will unfold in interesting ways and are worth paying attention to. This attitude will both enrich their own lives and make their partners' exploration easier.

Finally, they should understand their role as supporters. Psychologists call this role in a relationship the *secure base* and see it as vital to the other partner's growth. Originally identified and described by the psychologist John Bowlby, the secure base allows us to stretch ourselves by stepping outside our comfort zone while someone by our side soothes our anxieties about doing so. Without overly interfering, supporters should encourage their partners' exploration and reflection, even if it means moving away from the comfortable relationship they've already established.

Being a secure base for a partner presents its own trap, however:

Asymmetric support

In some couples one partner consistently supports the other without receiving support in return. That's what happened to Camille and Pierre. Pierre's experience in his former marriage, in which his wife gave up her career for his, made him determined to support Camille, and he initially stepped up to be a secure base for her. Their lives were so packed, however,

that Camille had trouble finding the energy to return the favor. The result was that her exploration and reflection became an impediment to Pierre's, creating a developmental and relationship deadlock. It is important to remember that acting as a secure base does not mean annihilating your own wishes, atoning for past selfishness, or being perfect. You can be a wonderful supporter for your partner while requesting support in return and taking time for yourself. In fact, that will most likely make you a far better (and less resentful) supporter.

In my research I found that couples who make it through their second transition are those in which the partners encourage each other to do this work—even if it means that one of them is exploring and providing support at the same time.

Once the exploring partner has had a chance to determine what he or she wants in a career, a life, or a relationship, the next step is to make it happen—as a couple. Couples need to renegotiate the roles they play in each other's lives. Take Matthew and James, another pair I spoke with, who had risen through the professional ranks in their 18 years together. When Matthew realized that he wanted to get off what he called the success train—on which he felt like a mere passenger—both he and James had to let go of their identity as a power couple and revisit the career-prioritization agreement they had forged during their first transition. Initially Matthew was reluctant to talk to James about his doubts, because he questioned whether

James would still love him if he changed direction. When they started discussing this, however, they realized that their identity as a power couple had trapped them in a dynamic in which both needed to succeed but neither could outshine the other. Acknowledging and renegotiating this unspoken arrangement allowed James to shoot for his first senior executive position and Matthew to transition into the nonprofit sector. The time and care they took to answer their existential questions and renegotiate the roles they played in each other's lives set them up for a renewed period of growth in their careers and in their relationship.

Transition 3: Loss and Opportunity

Attending her mother's funeral was one of the most difficult experiences of Norah's life. It was the culmination of two years of immense change for her and her husband, Jeremy, who were in their late fifties. The change began when their fathers unexpectedly passed away within five weeks of each other, and they became caregivers for Norah's ailing mother just as their children were leaving the nest and their own careers were in flux.

Jeremy is a digital visual artist. His studio's main projects were ending because a big client was moving on. Though he was sad, he had become confident enough to feel excited about whatever might come next. Norah had been working for the same small agricultural machinery business for 26 years; she

had once wanted to change careers but felt that she couldn't do so while Jeremy was relying on her for emotional and logistical support. Now she was being asked to take an early retirement deal. She felt thrown on the scrap heap despite her long commitment to the company. No career, no parents, no children to care for—who was she now? She felt disoriented and adrift.

The third transition is typically triggered by shifting roles later in life, which often create a profound sense of loss. Careers plateau or decline; bodies are no longer what they once were; children, if there are any, leave home. Sometimes one partner's career is going strong while the other's begins to ebb. Having raced through decades of career growth and child-rearing, couples wake up with someone who may have changed since the time they fell in love. They may both feel that way. These changes again raise fundamental questions of identity: Who am I now? Who do I want to be for the rest of my life?

Although loss usually triggers it, the third transition heralds opportunity. Chances for late-in-life reinvention abound, especially in today's world. Life expectancy is rising across the globe, and older couples may have several decades of reasonably good health and freedom from intensive parenting responsibilities. As careers and work become more flexible, especially for those with experience, people can engage in multiple activities more easily than previous generations could —combining advisory or consulting work with board service,

for example. Their activities often include giving back to the community, leaving some kind of legacy, mentoring younger generations, rediscovering passions of their youth, or dedicating themselves more to friendships.

Their task in the third transition is to again reinvent themselves—this time in a way that is both grounded in past accomplishments and optimistic about possibilities for the future. They must mourn the old, welcome the new, figure out how the two fit together, and adjust their life path to support who they want to become.

One thing that struck me when I spoke to couples in their third transition is that it's most powerful when partners reinvent themselves together—not just reflecting jointly, as in the other transitions, but actually taking on a new activity or project side by side. When one is curious about a partner's life and work as well as one's own, an immense capacity for mutual revitalization is unlocked. I met many couples who were charting new paths out of this transition that involved a merging of their work—launching a new business together, for example.

The third transition also has its traps:

Unfinished business

For better or for worse, earlier relational patterns, approaches, decisions, and assumptions will influence how a couple's third transition unfolds. I found that the most common challenge in

managing this transition was overcoming regret about perceived failures in the way the partners had "worked" as a couple—how they had prioritized their careers, or how each partner had supported the other's development (or not).

To move through the third transition, couples must acknowledge how they got where they are and commit to playing new roles for each other in the future. For example, Norah and Jeremy had become stuck in a pattern in which Norah was Jeremy's supporter. By recognizing this—and both their roles in cementing it—they were able to become more mutually supportive.

Narrow horizons

By the time a couple reaches the third transition, they will probably have suffered their fair share of disappointments and setbacks. They may be tired from years of taking care of others, or just from staying on the treadmill. As their roles shift and doubts about their identities grow, reinvention may be beyond consideration. In addition, because previous generations retired earlier, didn't live as long, and didn't have access to the gig economy, many couples lack role models for what reinvention can look like at this stage of life. If they don't deliberately broaden their horizons, they miss opportunities to discover themselves anew.

So couples must explore again. Even more than in the second transition, they need to flirt with multiple possibilities. Like

healthy children, who are curious about the world, themselves, and those around them, they can actively seek new experiences and experiment, avoid taking things for granted, and constantly ask “Why?” Most of us suppress our childhood curiosity as life progresses and responsibilities pile up. But it is vital to overcome the fear of leaving behind a cherished self and allow ambitions and priorities to diversify. Exploring at this stage is rejuvenating.

Shifts in people’s roles and identities offer a perfect excuse to question their current work, life, and loves. Many people associate exploring with looking for new options, which is surely important. But it’s also about questioning assumptions and approaches and asking, “Is this really how things need to be?”

Having rebalanced their support for each other, Norah and Jeremy could open up to new possibilities. Having earned financial security from their previous work, they sought reinvention not only in their careers but also in their wider roles in the world. Encouraging each other, they both transitioned to portfolio working lives. Jeremy became a freelance digital visual artist, took a part-time role teaching young art students at a local college, and dedicated more time to his passion of dinghy sailing. Norah retrained to be a counselor working with distressed families and began volunteering at a local agricultural museum. With these new opportunities and more time for each other and their friends,

they felt newfound satisfaction with their work and with their relationship.

The challenges couples face at each transition are different but linked. In their first transition, the partners accommodate to a major life event by negotiating the roles they will play in each other's lives. Over time those roles become constraining and spark the restlessness and questioning that lead to the second transition. To successfully navigate the third transition, couples must address regrets and developmental asymmetries left over from their first two transitions.

No one right path or solution exists for meeting these challenges. Although the 50/50 marriage—in which housework and child care are divided equally between the partners, and their careers are perfectly synced—may seem like a noble ideal, my research suggests that instead of obsessively trying to maintain an even "score," dual-career couples are better off being relentlessly curious, communicative, and proactive in making choices about combining their lives.

The Spouse Factor

by Jane Stevenson

Performing well as a high-level recruiter requires understanding what makes your candidates tick—and not just at work. That's especially true if I'm asking them to consider a job that requires relocation. In many cases I already know something about a candidate's family life—including the spouse's or partner's professional status, the ages of their kids (if any), and whether they have elderly parents living nearby. In cases where I don't know, I find a way to ask, "Is there anything in your family situation we should be sensitive to?" If there is, it's important to know early on, especially if these issues could become "blockers."

In my 34 years of experience, the most difficult factor to overcome when recruiting a candidate who has to relocate for a new job has been children, especially those in high school or with special needs. (This is often true whether the candidate is married or divorced; moving can be especially hard for someone who shares custody with an ex.) Spouses are the second most frequent reason a candidate will be reluctant to relocate, especially if he or she is part of a couple in which both are pursuing ambitious professional paths.

When I'm trying to recruit one member of a dual-career couple, it's important to fully understand the other's career, and also the city to which I hope to relocate them. For so-called trailing spouses, the most challenging careers are physician or

lawyer in private practice or owner of a business that isn't portable. People in these situations have often spent years building a client base and a local reputation, and it's difficult to reproduce those in a new region. The size of the destination city is also a factor. If the candidate's partner works in a traditional corporate function, he or she will have an easier time finding a similar (or better) position in a big market like Los Angeles or New York than in a smaller city. If the trailing spouse travels frequently for work, being near a major airport is also vital.

If the candidate's partner won't or can't move, I'll often ask whether the couple would consider a "commuter marriage." Companies today are increasingly willing to let high-performing leaders commute or work remotely. However, they are much more willing to allow an existing employee to do so, because they know the person's track record; the risk feels higher with a new employee.

Sometimes we have to think creatively. A few years ago a colleague and I were recruiting a female candidate who was based in Europe for a job in Asia. She had a long-term partner who had a great job and was unwilling to move to Asia. So we looked at the likely career path of the candidate (if she took the new role) and concluded that if she did a great job in Asia, she'd most likely be promoted to a position at headquarters in the United States, where her partner *was* willing to move. So the two commuted for a couple of years, and then the woman

I'd recruited did get a top job at headquarters; her partner moved to the United States, bringing them back together.

My work gives me a unique window onto how couples manage these situations, but my views are also informed by a research project I led at Korn Ferry on the careers of female CEOs. We interviewed 57 current or former CEOs about their paths to the corner office, and the most striking takeaway was the importance of strong spousal support for women who aspire to top jobs. When discussing the factors that led to their success, most of the women spontaneously brought up their husbands' support. About half the CEOs had spouses with substantial careers; managing their dual careers involved complex calendar negotiations, turn taking, weighing of career decisions, a willingness to relocate, and significant help from housekeepers, nannies, and so forth. About a third had spouses or partners who, by the time the women became CEOs, were assuming primary responsibility for home and children; some were househusbands, and others were retired or worked part-time.

Each of us has only so much energy to utilize, and dealing with a partner who isn't truly rooting for you professionally saps that energy, limiting your potential. A few of the CEOs we interviewed said they had previously had unsupportive husbands or partners but ultimately went on to connect with more-supportive ones. They speculated that they wouldn't have attained the top job if they hadn't received the support

they needed. (Most male CEOs I've worked with say the same thing.) My children are now 24 and 21, and I tell them very bluntly: Choosing a spouse may be one of the most important career decisions you'll ever make, because that person will be either a support or a hindrance to your professional ambitions. So choose wisely.

I empathize with couples who struggle with these issues, because I've faced them myself. My husband is a pathologist. We've been married 37 years. For roughly the first two decades, his career took priority. We moved several times to accommodate his medical school training, residency, fellowships, and stint with the U.S. Air Force. I believe a sturdy flower can bloom anywhere, so I tried to look at those moves as opportunities: When we moved from California to Philadelphia, my job search led me to executive recruiting. When we moved to Texas, where my then employer had no office, I opened a new one, which was great experience. As my career has evolved, we've made changes. My travel schedule is insane. In 2007, when our children were much younger, my husband left his hospital job and started consulting to have more flexibility and to be more available for the kids. Since then we've considered my career the priority.

Having experienced this push and pull, I recognize that it's typically a phone call from someone like me—followed by a great job offer—that causes a couple to rethink their coequal arrangement. Very often, the resulting conversation will focus

on the upside opportunity. It's natural for partners to compare the potential of their careers and decide to prioritize the one with the higher ROI. In the past that was typically the man's, but today it's frequently the woman's.

At such moments, many dual-career couples will decide that one career has to take a backseat, or that the lesser-earning partner will make a leap of faith and hope that he or she can find (or create) a great job in a new city. When couples face this prospect, I remind them that choosing to prioritize one partner's career doesn't mean it will be that way forever. Careers are long. The partner who's stepping back right now may be able to step forward in the future. I like to think I have credibility when I make this argument, because I've experienced that shift myself.

Originally published in September–October 2019. Reprint R1905B

Building the AI-Powered Organization

by Tim Fountaine, Brian McCarthy, and Tamim Saleh

ARTIFICIAL INTELLIGENCE IS reshaping business—though not at the blistering pace many assume. True, AI is now guiding decisions on everything from crop harvests to bank loans, and once pie-in-the-sky prospects such as totally automated customer service are on the horizon. The technologies that enable AI, like development platforms and vast processing power and data storage, are advancing rapidly and becoming increasingly affordable. The time seems ripe for companies to capitalize on AI. Indeed, we estimate that AI will add $13 trillion to the global economy over the next decade.

Yet, despite the promise of AI, many organizations' efforts with it are falling short. We've surveyed thousands of executives about how their companies use and organize for AI and advanced analytics, and our data shows that only 8% of firms engage in core practices that support widespread adoption. Most firms have run only ad hoc pilots or are applying AI in just a single business process.

Why the slow progress? At the highest level, it's a reflection of a failure to rewire the organization. In our surveys and our work with hundreds of clients, we've seen that AI initiatives face formidable cultural and organizational barriers. But we've also seen that leaders who at the outset take steps to break down those barriers can effectively capture AI's opportunities.

Making the Shift

One of the biggest mistakes leaders make is to view AI as a plug-and-play technology with immediate returns. Deciding to get a few projects up and running, they begin investing millions in data infrastructure, AI software tools, data expertise, and model development. Some of the pilots manage to eke out small gains in pockets of organizations. But then months or years pass without bringing the big wins executives expected. Firms struggle to move from the pilots to companywide programs—and from a focus on discrete business problems, such as improved customer segmentation, to big business challenges, like optimizing the entire customer journey.

Leaders also often think too narrowly about AI requirements. While cutting-edge technology and talent are certainly needed, it's equally important to align a company's culture, structure, and ways of working to support broad AI adoption. But at most businesses that aren't born digital, traditional mindsets and ways of working run counter to those needed for AI.

To scale up AI, companies must make three shifts:

From siloed work to interdisciplinary collaboration

AI has the biggest impact when it's developed by cross-functional teams with a mix of skills and perspectives. Having business and operational people work side by side with analytics experts will ensure that initiatives address broad organizational priorities, not just isolated business issues. Diverse teams can also think through the operational changes new applications may require—they're likelier to recognize, say, that the introduction of an algorithm that predicts maintenance needs should be accompanied by an overhaul of maintenance workflows. And when development teams involve end users in the design of applications, the chances of adoption increase dramatically.

Idea in Brief

The Problem

Many companies' efforts to scale up artificial intelligence fall short. That's because only 8% of firms are engaging in core practices that support widespread adoption.

The Solution

Cutting-edge technology and talent are not enough. Companies must break down organizational and cultural barriers that stand in AI's way.

The Leadership Imperatives

Leaders must convey the urgency of AI initiatives and their benefits for all; spend at least as much on adoption as on technology; organize AI work on the basis of the company's AI maturity, business complexity, and innovation pace; and invest in AI education for everyone.

From experience-based, leader-driven decision making to data-driven decision making at the front line

When AI is adopted broadly, employees up and down the hierarchy will augment their own judgment and intuition with algorithms' recommendations to arrive at better answers than either humans or machines could reach on their own. But for this approach to work, people at all levels have to trust the algorithms' suggestions and feel empowered to make decisions—and that means abandoning the traditional top-down approach. If employees have to consult a higher-up before taking action, that will inhibit the use of AI.

Decision processes shifted dramatically at one organization when it replaced a complex manual method for scheduling events with a new AI system. Historically, the firm's event planners had used colored tags, pins, and stickers to track conflicts, participants' preferences, and other considerations. They'd often relied on gut instinct and on input from senior managers, who also were operating on their instincts, to make decisions. The new system rapidly analyzed the vast range of scheduling permutations, using first one algorithm to distill hundreds of millions of options into millions of scenarios, and then another algorithm to boil down those millions into just hundreds, ranking the optimal schedules for each participant. Experienced human planners then applied their expertise to make final decisions supported by the data, without the need to get input from their leaders. The planners adopted the tool readily, trusting its output because they'd helped set its

parameters and constraints and knew that they themselves would make the final call.

From rigid and risk-averse to agile, experimental, and adaptable
Organizations must shed the mindset that an idea needs to be fully baked or a business tool must have every bell and whistle before it's deployed. On the first iteration, AI applications rarely have all their desired functionality. A test-and-learn mentality will reframe mistakes as a source of discoveries, reducing the fear of failure. Getting early user feedback and incorporating it into the next version will allow firms to correct minor issues before they become costly problems. Development will speed up, enabling small AI teams to create minimum viable products in a matter of weeks rather than months.

Such fundamental shifts don't come easily. They require leaders to prepare, motivate, and equip the workforce to make a change. But leaders must first be prepared themselves. We've seen failure after failure caused by the lack of a foundational understanding of AI among senior executives. (Further on, we'll discuss how analytics academies can help leaders acquire that understanding.)

Setting Up for Success

To get employees on board and smooth the way for successful AI launches, leaders should devote early attention to several tasks:

Explaining why

A compelling story helps organizations understand the urgency of change initiatives and how all will benefit from them. This is particularly critical with AI projects, because fear that AI will take away jobs increases employees' resistance to it.

Leaders have to provide a vision that rallies everyone around a common goal. Workers must understand why AI is important to the business and how they'll fit into a new, AI-oriented culture. In particular, they need reassurance that AI will enhance rather than diminish or even eliminate their roles. (Our research shows that the majority of workers will need to adapt to using AI rather than be replaced by AI.)

When a large retail conglomerate wanted to get its employees behind its AI strategy, management presented it as an existential imperative. Leaders described the threat that digital retailers posed and how AI could help fend it off by improving the firm's operational efficiency and responsiveness. By issuing a call to arms in a fight for survival, management underscored the critical role that employees had to play.

In sharing their vision, the company's leaders put a spotlight on workers who had piloted a new AI tool that helped them optimize stores' product assortments and increase revenue. That inspired other workers to imagine how AI could augment and elevate their performance.

Anticipating unique barriers to change

Some obstacles, such as workers' fear of becoming obsolete, are common across organizations. But a company's culture may also have distinctive characteristics that contribute to resistance. For example, if a company has relationship managers who pride themselves on being attuned to customer needs, they may reject the notion that a machine could have better ideas about what customers want and ignore an AI tool's tailored product recommendations. And managers in large organizations who believe their status is based on the number of people they oversee might object to the decentralized decision making or reduction in reports that AI could allow.

In other cases, siloed processes can inhibit the broad adoption of AI. Organizations that assign budgets by function or business unit may struggle to assemble interdisciplinary agile teams, for example.

Some solutions can be found by reviewing how past change initiatives overcame barriers. Others may involve aligning AI initiatives with the very cultural values that seem like obstacles. At one financial institution with a strong emphasis on relationship banking, for example, leaders highlighted AI's ability to enhance ties with customers. The bank created a booklet for relationship managers that showed how combining their expertise and skills with AI's tailored product recommendations could improve customers' experiences and increase revenue and profit. The AI adoption program also included a contest for sales conversions driven by using the new tool; the winners'

achievements were showcased in the CEO's monthly newsletter to employees.

A relatively new class of expert, analytics translators, can play a role in identifying roadblocks. These people bridge the data engineers and scientists from the technical realm with the people from the business realm—marketing, supply chain, manufacturing, risk personnel, and so on. Translators help ensure that the AI applications developed address business needs and that adoption goes smoothly. Early in the implementation process, they may survey end users, observe their habits, and study workflows to diagnose and fix problems.

Understanding the barriers to change can not only inform leaders about how to communicate with the workforce but also help them determine where to invest, what AI initiatives are most feasible, what training should be offered, what incentives may be necessary, and more.

Budgeting as much for integration and adoption as for technology (if not more)

In one of our surveys nearly 90% of the companies that had engaged in successful scaling practices had spent more than half of their analytics budgets on activities that drove adoption, such as workflow redesign, communication, and training. Only 23% of the remaining companies had committed similar resources.

Consider one telecom provider that was launching a new AI-driven customer-retention program in its call center. The company invested simultaneously in AI model development and

in helping the center's employees transition to the new approach. Instead of just reacting to calls canceling service, they would proactively reach out to customers at risk of defection, giving them AI-generated recommendations on new offers they'd be likely to accept. The employees got training and on-the-job coaching in the sales skills needed to close the business. Coaches and managers listened in on their calls, gave them individualized feedback, and continually updated the training materials and call scripts. Thanks to those coordinated efforts, the new program reduced customer attrition by 10%.

Balancing feasibility, time investment, and value

Pursuing initiatives that are unduly difficult to implement or require more than a year to launch can sabotage both current and future AI projects.

Organizations needn't focus solely on quick wins; they should develop a portfolio of initiatives with different time horizons. Automated processes that don't need human intervention, such as AI-assisted fraud detection, can deliver a return in months, while projects that require human involvement, such as AI-supported customer service, are likely to pay off over a longer period. Prioritization should be based on a long-term (typically three-year) view and take into consideration how several initiatives with different time lines could be combined to maximize value. For example, to achieve a view of customers detailed enough to allow AI to do microsegmentation, a company might need to set up a number of sales and marketing

initiatives. Some, such as targeted offers, might deliver value in a few months, while it might take 12 to 18 months for the entire suite of capabilities to achieve full impact.

An Asian Pacific retailer determined that an AI initiative to optimize floor space and inventory placement wouldn't yield its complete value unless the company refurbished all its stores, reallocating the space for each category of goods. After much debate, the firm's executives decided the project was important enough to future profitability to proceed—but not without splitting it in two. Part one produced an AI tool that gave store managers recommendations for a few incremental items that would sell well in their outlets. The tool provided only a small fraction of the total return anticipated, but the managers could get the new items into stores immediately, demonstrating the project's benefits and building enthusiasm for the multiyear journey ahead.

Organizing for Scale

There's a lot of debate about where AI and analytics capabilities should reside within organizations. Often leaders simply ask, "What organizational model works best?" and then, after hearing what succeeded at other companies, do one of three things: consolidate the majority of AI and analytics capabilities within a central "hub"; decentralize them and embed them mostly in the business units ("the spokes"); or distribute them across both, using a hybrid ("hub-and-spoke") model. We've found that none

of these models is always better than the others at getting AI up to scale; the right choice depends on a firm's individual situation.

Consider two large financial institutions we've worked with. One consolidated its AI and analytics teams in a central hub, with all analytics staff reporting to the chief data and analytics officer and being deployed to business units as needed. The second decentralized nearly all its analytics talent, having teams reside in and report to the business units. Both firms developed AI on a scale at the top of their industry; the second organization grew from 30 to 200 profitable AI initiatives in just two years. And both selected their model after taking into account their organizations' structure, capabilities, strategy, and unique characteristics.

The hub

A small handful of responsibilities are always best handled by a hub and led by the chief analytics or chief data officer. These include data governance, AI recruiting and training strategy, and work with third-party providers of data and AI services and software. Hubs should nurture AI talent, create communities where AI experts can share best practices, and lay out processes for AI development across the organization. Our research shows that companies that have implemented AI on a large scale are three times as likely as their peers to have a hub and 2.5 times as likely to have a clear methodology for creating models, interpreting insights, and deploying new AI capabilities.

Hubs should also be responsible for systems and standards related to AI. These should be driven by the needs of a firm's initiatives, which means they should be developed gradually, rather than set up in one fell swoop, before business cases have been determined. We've seen many organizations squander significant time and money—spending hundreds of millions of dollars—up front on companywide data-cleaning and data-integration projects, only to abort those efforts midway, realizing little or no benefits.

In contrast, when a European bank found that conflicting data-management strategies were hindering its development of new AI tools, it took a slower approach, making a plan to unify its data architecture and management over the next four years as it built various business cases for its AI transformation. This multiphase program, which also includes an organizational redesign and a revised talent strategy, is expected to have an annual impact of more than $900 million.

The spokes

Another handful of responsibilities should almost always be owned by the spokes, because they're closest to those who will be using the AI systems. Among them are tasks related to adoption, including end-user training, workflow redesign, incentive programs, performance management, and impact tracking.

To encourage customers to embrace the AI-enabled services offered with its smart, connected equipment, one manufacturer's

sales and service organization created a “SWAT team” that supported customers using the product and developed a pricing plan to boost adoption. Such work is clearly the bailiwick of a spoke and can’t be delegated to an analytics hub.

The gray area

Much of the work in successful AI transformations falls into a gray area in terms of responsibility. Key tasks—setting the direction for AI projects, analyzing the problems they’ll solve, building the algorithms, designing the tools, testing them with end users, managing the change, and creating the supporting IT infrastructure—can be owned by either the hub or the spoke, shared by both, or shared with IT. (See the exhibit “Organizing AI for scale.”) Deciding where responsibility should lie within an organization is not an exact science, but it should be influenced by three factors:

Organizing AI for scale

AI-enabled companies divide key roles between a hub and spokes. A few tasks are always owned by the hub, and the spokes always own execution. The rest of the work falls into a gray area, and a firm’s individual characteristics determine where it should be done.

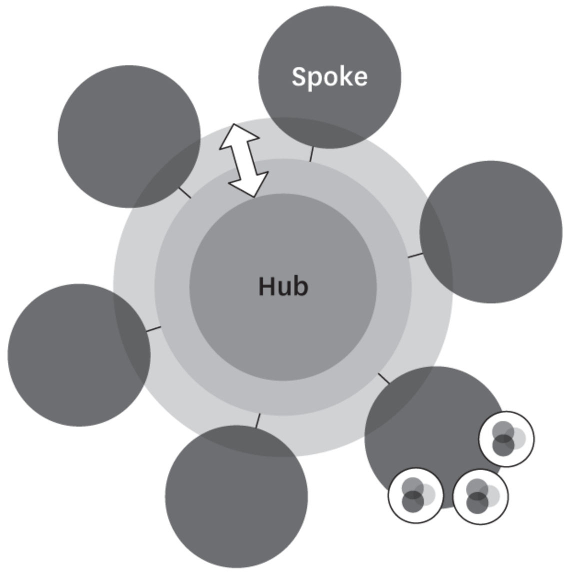

Governing coalition

A team of business, IT, and analytics leaders that share accountability for the AI transformation

Hub

A central group headed by a C-level analytics executive who aligns strategy

Responsibilities

- Talent recruitment and training strategy
- Performance management
- Partnerships with providers of data and AI services and software
- AI standards, processes, policies

Gray area

Work that could be owned by the hub or spokes or shared with IT

Responsibilities

- Project direction, delivery, change management
- Data strategy, data architecture, code development
- User experience
- IT infrastructure
- Organizational capability assessment, strategy, funding

Spoke

A business unit, function, or geography, which assigns a manager to be the AI product owner and a business analyst to assist him or her

Responsibilities

- Oversight of execution teams
- Solution adoption
- Performance tracking

Execution teams

Assembled from the hub, spoke, and gray area for the duration of the project

Key roles

- Product owner
- Analytics translator
- Data scientist
- Data engineer
- Data architect
- Visualization specialist
- UI designer
- Business analyst

The maturity of AI capabilities. When a company is early in its AI journey, it often makes sense for analytics executives, data scientists, data engineers, user interface designers, visualization specialists who graphically interpret analytics findings, and the like to sit within a hub and be deployed as needed to the spokes. Working together, these players can establish the company's core AI assets and capabilities, such as common analytics tools, data processes, and delivery methodologies. But as time passes and processes become standardized, these experts can reside within the spokes just as (or more) effectively.

Business model complexity. The greater the number of business functions, lines of business, or geographies AI tools will support, the greater the need to build guilds of AI experts (of, say, data scientists or designers). Companies with complex businesses often consolidate these guilds in the hub and then assign them out as needed to business units, functions, or geographies.

The pace and level of technical innovation required. When they need to innovate rapidly, some companies put more gray-area strategy and capability building in the hub so they can monitor industry and technology changes better and quickly deploy AI resources to head off competitive challenges.

Let's return to the two financial institutions we discussed earlier. Both faced competitive pressures that required rapid

innovation. However, their analytics maturity and business complexity differed.

The institution that placed its analytics teams within its hub had a much more complex business model and relatively low AI maturity. Its existing AI expertise was primarily in risk management. By concentrating its data scientists, engineers, and many other gray-area experts within the hub, the company ensured that all business units and functions could rapidly access essential know-how when needed.

The second financial institution had a much simpler business model that involved specializing in fewer financial services. This bank also had substantial AI experience and expertise. So it was able to decentralize its AI talent, embedding many of its gray-area analytics, strategy, and technology experts within the business-unit spokes.

As these examples suggest, some art is involved in deciding where responsibilities should live. Every organization has distinctive capabilities and competitive pressures, and the three key factors must be considered in totality, rather than individually. For example, an organization might have high business complexity and need very rapid innovation (suggesting it should shift more responsibilities to the hub) but also have very mature AI capabilities (suggesting it should move them to the spokes). Its leaders would have to weigh the relative importance of all three factors to determine where, on balance, talent would most effectively be deployed. Talent levels (an

element of AI maturity) often have an outsize influence on the decision. Does the organization have enough data experts that, if it moved them permanently to the spokes, it could still fill the needs of all business units, functions, and geographies? If not, it would probably be better to house them in the hub and share them throughout the organization.

Oversight and execution

While the distribution of AI and analytics responsibilities varies from one organization to the next, those that scale up AI have two things in common:

A governing coalition of business, IT, and analytics leaders. Fully integrating AI is a long journey. Creating a joint task force to oversee it will ensure that the three functions collaborate and share accountability, regardless of how roles and responsibilities are divided. This group, which is often convened by the chief analytics officer, can also be instrumental in building momentum for AI initiatives, especially early on.

Assignment-based execution teams. Organizations that scale up AI are twice as likely to set up interdisciplinary teams within the spokes. Such teams bring a diversity of perspectives together and solicit input from frontline staff as they build, deploy, and monitor new AI capabilities. The teams are usually assembled at the outset of each initiative and draw skills from both the hub and the spokes. Each generally includes the manager in charge of the new AI tool's success (the "product owner"), translators, data

architects, engineers and scientists, designers, visualization specialists, and business analysts. These teams address implementation issues early and extract value faster.

For example, at the Asian Pacific retailer that was using AI to optimize store space and inventory placement, an interdisciplinary execution team helped break down walls between merchandisers (who determined how items would be displayed in stores) and buyers (who chose the range of products). Previously, each group had worked independently, with the buyers altering the AI recommendations as they saw fit. That led to a mismatch between inventory purchased and space available. By inviting both groups to collaborate on the further development of the AI tool, the team created a more effective model that provided a range of weighted options to the buyers, who could then choose the best ones with input from the merchandisers. At the end of the process, gross margins on each product category that had applied the tool increased by 4% to 7%.

Educating Everyone

To ensure the adoption of AI, companies need to educate everyone, from the top leaders down. To this end some are launching internal AI academies, which typically incorporate classroom work (online or in person), workshops, on-the-job training, and even site visits to experienced industry peers. Most academies initially hire external faculty to write the curricula

and deliver training, but they also usually put in place processes to build in-house capabilities.

Every academy is different, but most offer four broad types of instruction:

Leadership

Most academies strive to give senior executives and business-unit leaders a high-level understanding of how AI works and ways to identify and prioritize AI opportunities. They also provide discussions of the impact on workers' roles, barriers to adoption, and talent development, and offer guidance on instilling the underlying cultural changes required.

10 Ways to Derail an AI Program

DESPITE BIG INVESTMENTS, many organizations get disappointing results from their AI and analytics efforts. What makes programs go off track? Companies set themselves up to fail when:

1. They lack a clear understanding of advanced analytics, staffing up with data scientists, engineers, and other key players without realizing how advanced and traditional analytics differ.
2. They don't assess feasibility, business value, and time horizons, and launch pilots without thinking through how to balance short-term wins in the first year with longer-term payoffs.
3. They have no strategy beyond a few use cases, tackling AI in an ad hoc way without considering the big-picture opportunities and

threats AI presents in their industry.

4. They don't clearly define key roles, because they don't understand the tapestry of skill sets and tasks that a strong AI program requires.
5. They lack "translators," or experts who can bridge the business and analytics realms by identifying high-value use cases, communicating business needs to tech experts, and generating buy-in with business users.
6. They isolate analytics from the business, rigidly centralizing it or locking it in poorly coordinated silos, rather than organizing it in ways that allow analytics and business experts to work closely together.
7. They squander time and money on enterprisewide data cleaning instead of aligning data consolidation and cleanup with their most valuable use cases.
8. They fully build out analytics platforms before identifying business cases, setting up architectures like data lakes without knowing what they'll be needed for and often integrating platforms with legacy systems unnecessarily.
9. They neglect to quantify analytics' bottom-line impact, lacking a performance management framework with clear metrics for tracking each initiative.
10. They fail to focus on ethical, social, and regulatory implications, leaving themselves vulnerable to potential missteps when it comes

to data acquisition and use, algorithmic bias, and other risks, and exposing themselves to social and legal consequences.

For more details, read "Ten Red Flags Signaling Your Analytics Program Will Fail" on McKinsey.com.

Analytics
Here the focus is on constantly sharpening the hard and soft skills of data scientists, engineers, architects, and other employees who are responsible for data analytics, data governance, and building the AI solutions.

Translator
Analytics translators often come from the business staff and need fundamental technical training—for instance, in how to apply analytical approaches to business problems and develop AI use cases. Their instruction may include online tutorials, hands-on experience shadowing veteran translators, and a final "exam" in which they must successfully implement an AI initiative.

End user
Frontline workers may need only a general introduction to new AI tools, followed by on-the-job training and coaching in how to use them. Strategic decision makers, such as marketers and finance staff, may require higher-level training sessions that incorporate real business scenarios in which new tools improve decisions about, say, product launches.

Reinforcing the Change

Most AI transformations take 18 to 36 months to complete, with some taking as long as five years. To prevent them from losing momentum, leaders need to do four things:

Walk the talk

Role modeling is essential. For starters, leaders can demonstrate their commitment to AI by attending academy training.

But they also must actively encourage new ways of working. AI requires experimentation, and often early iterations don't work out as planned. When that happens, leaders should highlight what was learned from the pilots. That will help encourage appropriate risk taking.

The most effective role models we've seen are humble. They ask questions and reinforce the value of diverse perspectives. They regularly meet with staff to discuss the data, asking questions such as "How often are we right?" and "What data do we have to support today's decision?"

The CEO of one specialty retailer we know is a good example. At every meeting she goes to, she invites attendees to share their experience and opinions—and offers hers last. She also makes time to meet with business and analytics employees every few weeks to see what they've done—whether it's launching a new pilot or scaling up an existing one.

Make businesses accountable

It’s not uncommon to see analytics staff made the owners of AI products. However, because analytics are simply a means of solving business problems, it’s the business units that must lead projects and be responsible for their success. Ownership ought to be assigned to someone from the relevant business, who should map out roles and guide a project from start to finish. Sometimes organizations assign different owners at different points in the development life cycle (for instance, for proof of value, deployment, and scaling). That’s a mistake too, because it can result in loose ends or missed opportunities.

A scorecard that captures project performance metrics for all stakeholders is an excellent way to align the goals of analytics and business teams. One airline company, for instance, used a shared scorecard to measure rate of adoption, speed to full capability, and business outcomes for an AI solution that optimized pricing and booking.

Track and facilitate adoption

Comparing the results of decisions made with and without AI can encourage employees to use it. For example, at one commodity company, traders learned that their non-AI-supported forecasts were typically right only half the time—no better than guessing. That discovery made them more open to AI tools for improved forecasting.

Teams that monitor implementation can correct course as needed. At one North American retailer, an AI project owner saw store managers struggling to incorporate a pilot’s output into

their tracking of store performance results. The AI's user interface was difficult to navigate, and the AI insights generated weren't integrated into the dashboards the managers relied on every day to make decisions. To fix the issue, the AI team simplified the interface and reconfigured the output so that the new data stream appeared in the dashboard.

Provide incentives for change

Acknowledgment inspires employees for the long haul. The CEO of the specialty retailer starts meetings by shining a spotlight on an employee (such as a product manager, a data scientist, or a frontline worker) who has helped make the company's AI program a success. At the large retail conglomerate, the CEO created new roles for top performers who participated in the AI transformation. For instance, he promoted the category manager who helped test the optimization solution during its pilot to lead its rollout across stores—visibly demonstrating the career impact that embracing AI could have.

Finally, firms have to check that employees' incentives are truly aligned with AI use. This was not the case at a brick-and-mortar retailer that had developed an AI model to optimize discount pricing so that it could clear out old stock. The model revealed that sometimes it was more profitable to dispose of old stock than to sell it at a discount, but the store personnel had incentives to sell everything, even at steep discounts. Because the AI recommendations contradicted their standard, rewarded practice, employees became suspicious of the tool and ignored it.

Since their sales incentives were also closely tied to contracts and couldn't easily be changed, the organization ultimately updated the AI model to recognize the trade-off between profits and the incentives, which helped drive user adoption and lifted the bottom line.

The actions that promote scale in AI create a virtuous circle. The move from functional to interdisciplinary teams initially brings together the diverse skills and perspectives and the user input needed to build effective tools. In time, workers across the organization absorb new collaborative practices. As they work more closely with colleagues in other functions and geographies, employees begin to think bigger—they move from trying to solve discrete problems to completely reimagining business and operating models. The speed of innovation picks up as the rest of the organization begins to adopt the test-and-learn approaches that successfully propelled the pilots.

As AI tools spread throughout the organization, those closest to the action become increasingly able to make decisions once made by those above them, flattening organizational hierarchies. That encourages further collaboration and even bigger thinking.

The ways AI can be used to augment decision making keep expanding. New applications will create fundamental and sometimes difficult changes in workflows, roles, and culture, which leaders will need to shepherd their organizations through

carefully. Companies that excel at implementing AI throughout the organization will find themselves at a great advantage in a world where humans and machines working together outperform either humans or machines working on their own.

Originally published in July–August 2019. Reprint R1904C

Leading a New Era of Climate Action

by Andrew Winston

CLIMATE CHANGE IS A GLOBAL EMERGENCY. It's threatening crops, water supplies, infrastructure, and livelihoods. It's damaging the broader economy and company bottom lines *today*, not in some distant future. In recent years AT&T has spent $874 million on repairs after natural disasters that the company ties to climate change. The reinsurance leader Swiss Re has seen large increases in payouts for damage caused by extreme weather events—$2.5 billion more in 2017 than it had predicted—a trend that CEO Christian Mumenthaler attributes to rising global temperatures. If we don't move quickly toward action on climate, says Mark Carney, the Bank of England governor, we'll see company bankruptcies and raise the odds of systemic economic collapse.

Corporate leaders are at last absorbing this; nearly every large company has significant plans to cut carbon emissions and is acting. But given the scale of the crisis and the pace at which it's developing, these efforts are woefully inadequate. Critical UN reports in 2018 and 2019 make two things clear: (1) To avoid *some* of the worst outcomes of climate change, the world must cut

carbon emissions by 45% by 2030 and eliminate them entirely by midcentury. (2) Current government plans and commitments are not remotely close to putting us on that path. Emissions are still rising.

Countries, cities, and businesses need to move simultaneously along two paths: reducing emissions dramatically (mitigation) and investing in resilience while planning for vast change (adaptation). My focus here is on mitigation, because adaptation alone—building ever-higher walls to keep out the sea and simply turning up the air-conditioning as the outdoors becomes uninhabitable—won't save us. If we allow climate change to destroy the plant and animal ecosystems we rely on, there will be no replacements. The good news is that business has enormous potential to profitably cut emissions faster and even more.

If the main question for business were still "Which actions will both cut emissions and create short-term value?" we know the answer: slash carbon in energy-intensive industries and in operations, transportation, and buildings; buy lots of renewable energy, which is strategically smart because it has been competitive with fossil fuels for years; reduce waste, particularly in critical sectors such as food and agriculture; expand the use of circular business models that minimize resource use; embed climate change metrics in corporate systems and key performance indicators; and more. Again, most companies have begun to take advantage of these "basic" opportunities and will

accelerate adoption as they see the payoff grow. So let's assume that they will continue down this path. Then what?

Given the urgency, we must ask a different, and harder, question: "What are *all* the things business can possibly do with its vast resources?" What capital—financial, human, brand, and political—can companies bring to bear?

Drawing on 20 years of consulting to global corporations and working on climate change issues, I see three actions that companies must now focus on to drive deeper change:

- using political influence to demand aggressive climate policies around the world
- empowering suppliers, customers, and employees to drive change
- rethinking investments and business models to eliminate waste and carbon throughout the economy

Idea in Brief

Climate change is a global emergency that threatens crops, water supplies, infrastructure, and livelihoods. It's damaging the economy and company bottom lines. Most large companies are cutting carbon emissions, but given the scale of the crisis, these efforts are sadly inadequate. Companies need to mobilize, says Andrew Winston, to deal with this unprecedented global problem. He draws on 20 years of consulting for global corporations to recommend three actions:

- Use political influence to demand aggressive climate policies.
- Empower suppliers, customers, and employees to drive change.

- Rethink investments and business models to eliminate waste and carbon.

These actions may feel unnatural to some executives if they appear to put larger interests ahead of immediate shareholder profits. But the tide is turning on the very idea of shareholder primacy. The roughly 200 largest multinationals based in the United States recently declared, through the Business Roundtable, that they will no longer focus solely on shareholders or on the short run. We are at a pivotal moment as the climate crisis propels companies' growing sense of social purpose. The result, I believe, is the will needed to finally achieve this deeper change.

What's in It for Us?

Before I dig into the three areas of change, it's fair to ask why a company would commit to such challenging and possibly risky initiatives. One argument is macro/societal and the other is microeconomic. The former is straightforward: Companies need healthy people and a viable planet; with expensive runaway climate change on the horizon, they have an economic imperative and a moral responsibility to do everything they can to ensure a thriving world. As Unilever's former CEO Paul Polman says, "Business simply can't be a bystander in a system that gives it life in the first place." And let's not forget that even as they

pursue their own self-interest, executives sometimes just do what they believe is the right thing, which may or may not pay off—from ceasing to sell assault weapons at Dick's Sporting Goods and Walmart to funding by Apple and Microsoft of programs to reduce homelessness in their neighborhoods.

The Big Idea: Mobilizing on Climate

"Leading a New Era of Climate Action" is the lead article of HBR's **The Big Idea: Mobilizing on Climate.** Read the rest of the series at hbr.org/climate:

- "Tough Business Questions about the Climate Crisis," by Andy Robinson
- "What Do People Really Believe about Climate Change?" by Gretchen Gavett
- "Your Company's Next Leader on Climate Is … the CFO," by Laura Palmeiro and Delphine Gibassier
- "The New Business of Garbage," by Laura Amico
- "A Better Way to Talk about the Climate Crisis," by Gretchen Gavett
- "Is Your Trade Group Blocking Climate Action?" by Sheldon Whitehouse

The microeconomic argument, however, is often overlooked. Stakeholders, particularly customers and employees, have

increasingly high standards for the companies they buy from and work for. Business customers are demanding more sustainability performance from suppliers every year. Consumers are seeking out sustainable brands (50% of consumer packaged goods growth from 2013 to 2018 came from sustainability-marketed products), and Deloitte's global surveys show that up to 87% of the under-40 crowd—the Millennials who will make up 75% of the global workforce in five years—believes that a company's success should be measured in more than just financial terms. And nine in 10 members of Gen Z agree that companies have a responsibility to engage with environmental and social issues.

Employees are now directly pressuring their companies to do more on climate, particularly in the tech sector. In direct and public appeals, Google employees have asked their executives to cut ties to climate deniers, and Microsoft's employees staged a walkout in protest of the company's "complicity in the climate crisis." At Amazon more than 8,700 workers have signed an open letter to CEO Jeff Bezos with a list of demands, including developing a plan to get to zero emissions and eliminating donations to climate-denying legislators. Their efforts clearly played a part in pushing Bezos to announce large ambitions to be carbon neutral by 2040 and to buy 100,000 electric vehicles.

Alarming forecast: current climate policies are grossly inadequate

To hold global warming to 1.5° Celsius above preindustrial levels and prevent the worst impacts of climate change, the world must cut carbon

emissions to zero by midcentury. Yet emissions are still rising, and under existing policies reductions won't begin to approach what's needed. If we stay on the current path, temperatures will probably increase by about 3° C, with catastrophic effects.

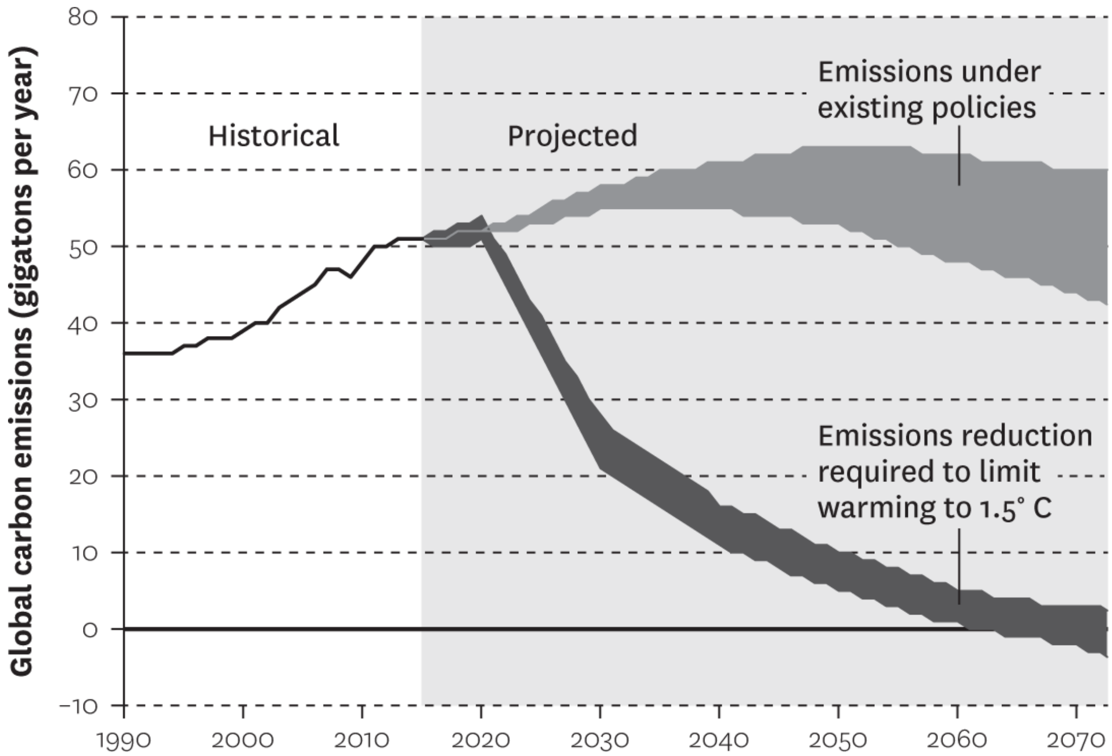

Source: Climate Action Tracker.

Note: Bandwidths represent high and low emissions estimates.

Because of pressure like this, along with increasingly dire warnings from climate scientists and global bodies including the UN, corporate efforts to reduce emissions have become table stakes—something any company *must* do to earn respect from employees and customers. And what is common and accepted practice, regardless of the short-term ROI, can sometimes shift very quickly. Consider that nobody could prove the value of

diversity and inclusion when companies first dove into that issue. Now we have good data—but the norms changed first.

I've seen firsthand how this can play out on sustainability issues. Nearly six years ago, in my book *The Big Pivot*, I advocated setting science-based emissions-reduction goals. Virtually no companies were doing that then, and I argued with many who wondered why a company would set a goal not required by law. Now, owing to peer pressure—and because it's rational—those goals are all but standard for big companies, with about 750 signed up and more than 200 committing to 100% renewable energy. They moved from "Why would we?" to "You're a laggard if you don't."

The first companies to try the most innovative sustainability strategies are generally B Corps or purpose-driven, privately held businesses like Patagonia and IKEA, which have more leeway to experiment. The story is similar for many of the next-gen climate ideas I lay out below: Big public companies are just dipping their toes in the water, while smaller, nimbler, sustainability-focused companies take the lead. Their examples matter, because over the past decade the largest firms started emulating the midsize leaders—or just buying them. To mitigate the worst effects of climate change, more companies need to follow, and fast.

Let's return now to the three broad activities that every company, big or small, must undertake.

1. Use Political Influence for Climate Good

Given the scale of the climate crisis, business alone can't solve it. But business does have a powerful tool beyond its own practices and products: extensive and deep tendrils in the halls of political power. All over the world, but especially in market economies, companies have enormous influence over governments and politicians. Through large campaign donations and—in the United States after the Supreme Court case *Citizens United*—nearly unlimited spending on political ads, the corporate agenda gets an outsize voice in society. How can and should companies use that power?

Business's government relations have traditionally been aimed at reshaping or fighting regulations. But over the past few years many companies have, at least on the surface, been supporting some climate policy. Hundreds of multinationals with operations in the U.S. have signed statements such as "We Are Still In" and the recent "United for the Paris Agreement" to let the world know that they will cut emissions in keeping with the Paris Climate Accords and that they want the U.S. government to stay aboard, despite announcements that it would not. Another group of large companies called for the world to hold warming to just 1.5 degrees Celsius. Signatories came from every corner of the planet: Sweden (Electrolux), Japan (ASICS), India (Mahindra Group), Switzerland (Nestlé), Germany (SAP), and many other places and sectors.

But statements alone are inadequate. Companies must lobby for the policies that will lead to a low-carbon future, and senior

executives need to show up in person. Without collective government action, we have little chance of avoiding the direst outcomes of climate change. One industry—fossil fuels—has had a dominant, decades-long influence on climate policies in world capitals, and for good reason: Policies aimed at reducing emissions pose an existential threat to the business. Companies in every other sector must grasp that climate change, which may spin out of control without enlightened policies, is an existential threat to *their* businesses.

For the most part, nonfossil fuel companies engage only in occasional special lobbying days organized by the likes of Ceres, the American Sustainable Business Council, and Business Climate Leaders. Those events are important, of course, but even the groups themselves acknowledge that the number of big companies with a consistent climate-action focus is small. As Joe Britton, a former chief of staff for U.S. senator Martin Heinrich, told me, these temporary “fly-ins” are better than nothing, but they are overshadowed by the daily swarm of fossil fuel lobbyists. In response, Britton left his position to create a new lobbying organization, with the help of other Capitol Hill insiders, to deploy a fuller and more constant political message to Congress on climate.

There’s also a major disconnect between what companies say about their commitments to fight climate change and what those who represent them—the trade associations or even their own government relations people—actually push for. As transparency

increases, companies should worry about any gap between their sustainability commitments and their lobbying. An NGO, Australia's LobbyWatch, is calling out the mining giant BHP and others for such disconnects. And the UK-based influencemap.org is tracking corporate lobbying activity on climate at hundreds of companies and publicly highlighting hypocrisy.

For leaders, aggressive climate lobbying is not just about appearances; it can create advantage. If 100% of your energy comes from renewables, a price on carbon won't affect your own cost structure much. And if you make products or provide services that help reduce emissions, you benefit from tighter carbon controls. That's surely one reason that Germany's Siemens, with a portfolio of products that improve energy efficiency, states that its top political engagement goal is "combating Climate Change."

Hugh Welsh, the president for North America at DSM, a large Dutch company that offers nutrition, health, and sustainable-living products and solutions, can attest to this. He has worked for years to bring a business voice on climate to the halls of political power. Welsh says he does this for two reasons: principles and pragmatism. About the former, he says, "Over 10 years as president, I've developed political capital. I can use that just for strategic things for the business, but I can also use that to improve the world." About the latter, he notes that DSM serves several sustainability-focused product markets, so a proactive role on sustainability and climate policy fits its strategy.

When Welsh makes the case to skeptical executives, leaders, and trade groups—such as the recalcitrant U.S. Chamber of Commerce, with which he worked for two years to flip its position on climate—he says, “If you don’t evolve your position, you’ll be on the wrong side of history … your partners and customers will leave in droves.”

Rising temperatures, rising risks: flooding cities

If the global temperature were to increase by …

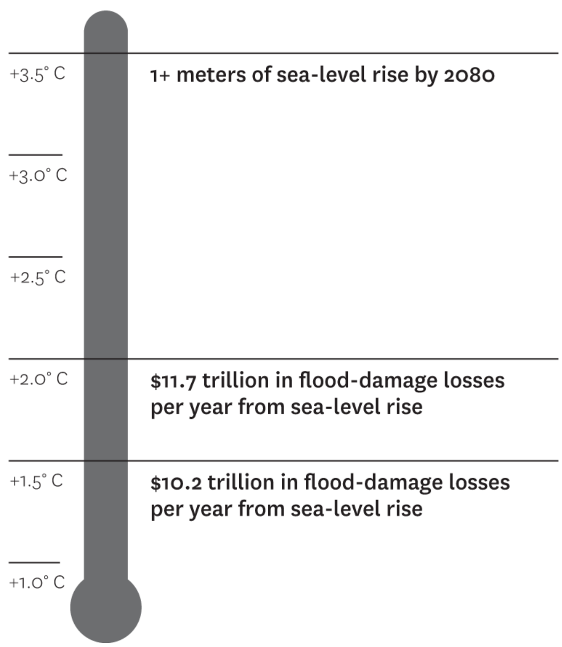

Source: World Resources Institute.

So what policies should companies advocate? To move the world to a low-carbon future, we need bold plans in a few key areas: pricing carbon and mobilizing capital to shift to low-carbon systems; rapidly raising performance standards and phasing out

old technologies for big energy users like cars and buildings; and enabling transparency and efforts to reduce human suffering.

These priorities apply in most geographies, but of course policy formation and the relationship between business and government vary widely across countries. Approaches in command-and-control economies must vary from those in sprawling capitalist systems.

Policies may take years to have an effect, so these efforts must be made soon. It's time for companies to use their substantial political influence to proactively support laws that make high-carbon products and choices more expensive, mobilize capital toward a clean economy, support systems change, and help deal with adaptation and the human costs of shifts to clean technology.

Climate Policies Companies Should Fight For

A long list of possible government policies could create the conditions for rapid emissions reductions. But the following are probably the most important for business to get behind. These will fix market failures, shift capital toward low-carbon investments, and set a high bar for low-carbon products.

Implement a rapidly rising price on carbon, coupled with massive shifts in subsidies from fossil fuels to clean tech and low-carbon production methods.

Create incentives for farmers to move from industrial to regenerative agriculture.

Fund increased material capture (recycling, reuse, repair) to encourage a circular economy.

Mobilize capital and R&D that pulls public and private investment into cleaner tech. For example, the Danish aviation sector has proposed a climate tax on all flights from Denmark, earmarked for a fund to research green solutions and climate-neutral fuels.

Introduce high performance standards for the big energy users, including cars, buildings, and HVAC systems.

Encourage phaseouts and phase-ins such as by mandating low-global-warming-potential refrigerants and net-zero buildings with renewables and banning gas-guzzlers. Some countries have set a date for stopping the sale of internal combustion engines: Norway by 2025, Sweden and Denmark by 2030, and France and Sri Lanka by 2040.

Prioritize transparency through, for example, the Task Force on Climate-related Financial Disclosures, which provides guidelines for companies reporting their material risks from climate change, and product labels with carbon-footprint information, much like the calorie and nutrition counts on food labels.

Fund resources for adaptation, such as resilience planning in cities, the relocation of citizens, and retraining for those from older

sectors that will rapidly decline.

---

2. Leverage Stakeholder Relationships

At the same time, companies should wield their other superpower: vast influence over value chain partners and deep connections to their customers and employees. Big consumer products companies like P&G and Unilever often rightly brag that they serve billions of people every day. More than 275 million people visit a Walmart every week. Companies employ hundreds of millions of us. And with nearly $33 trillion in revenues across the *Fortune* Global 500 alone, it's safe to assume that many trillions go to suppliers. Imagine if companies used those touch points, their buying power, and all their communications and advertising clout to catalyze change across business and society.

Suppliers

In recent years corporations have ratcheted up the pressure on their suppliers to operate more sustainably. Big buyers increasingly want to see progress—backed up by data—in a supplier's carbon footprint, resource use, human rights and labor performance, and much more. General Mills, Kellogg, IKEA, and Hewlett Packard Enterprise have all set science-based carbon goals for their suppliers. Others, including GSK, H&M, Toyota, and Schneider Electric, have committed to carbon neutrality or negativity (eliminating more carbon than is produced) in their entire value chains by 2040 or 2050.

Commitments like these are becoming the norm. But what else is possible? What are boundary-pushing companies doing to drive change? I see future supply-chain climate leadership in three key areas: providing capital, driving innovation and collaboration, and using purchasing power to choose suppliers on the basis of emissions performance.

Financial assistance and capital. Making a business more sustainable is profitable, but it may still require investments and capital. Companies that ask suppliers to change how they do business can help, especially with smaller players. For example, in mid-2018, after achieving 100% renewable energy in its own operations, Apple launched the China Clean Energy Fund, a joint pool of $300 million to help suppliers buy one gigawatt of renewable energy, and the fund's first big wind farms went up last year. Similarly, IKEA recently committed €100 million to help first-tier suppliers make the shift. In another innovative approach, an industrial company I work with, Ingersoll Rand (better known by its brands Thermo King and Trane), financed a large renewable energy project and then invited suppliers to offset their emissions by buying portions of the energy production. And beyond encouraging renewables, some leaders, such as Levi's and Walmart, have worked with HSBC and other banks to provide lower interest rates to suppliers that score well on sustainability performance.

Rising temperatures, rising risks: food shortages

If the global temperature were to increase by ...

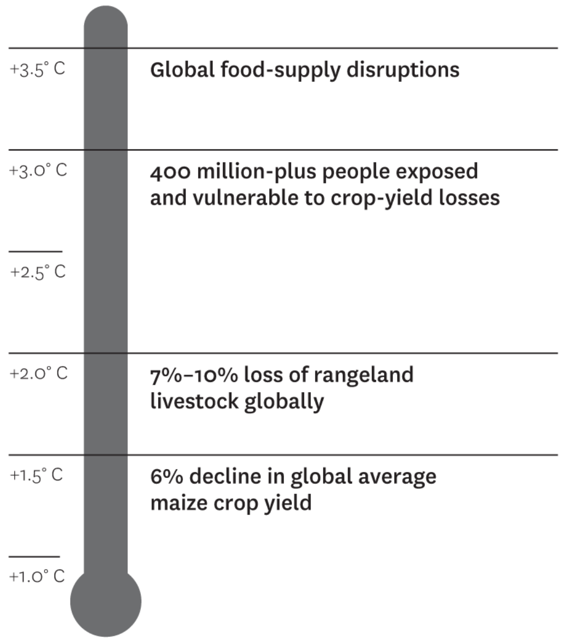

Source: World Resources Institute.

Joint innovation. I also recently watched the head of procurement at Ingersoll Rand tell hundreds of suppliers that his company would no longer choose vendors on the basis of pricing and quality alone. Now, he said, suppliers would need to innovate *with* the company to make its products more energy- and carbon-efficient. This is a great way to drive value chain innovation, but sectorwide collaboration can have an even bigger impact.

Consider that Walmart and Target, which are traditionally competitors, worked together with the NGO Forum for the Future (on whose board I serve) to create the Beauty and Personal Care Sustainability Project—a creative attempt at improving the environmental and social footprint of all the

products we put on our bodies. They brought together big CPG companies such as P&G and Unilever and their chemical suppliers to rethink ingredients, packaging, and more to reduce health and environmental impacts. Apple has dived deep into its supply chain to make its ubiquitous tech products lower-carbon, including through a joint venture with Rio Tinto and Alcoa to develop and commercialize an aluminum-smelting process with vastly lower greenhouse gas emissions and lower costs.

Purchasing power. For years many companies have agreed to work with lagging suppliers to improve their sustainability performance. But the world can no longer afford to wait for slow adopters. Companies should cut them loose and shift their purchasing dollars toward the low-carbon leaders—which are often the best-run suppliers anyway. VF Corporation, the home of brands such as Vans and The North Face, stopped buying leather from Brazil because government policy there was encouraging Amazon rain forest destruction.

Retailers should make carbon performance a buying priority. Mainstream mega-retailers like Walmart and Target have pressured suppliers for years to make their offerings more sustainable, but they could do much more to support those that are best at reducing emissions in their operations or through their products. They could, for example, permanently (not just on Earth Day) devote endcaps or special promotion areas—their highest-value real estate—to drive business to the lowest-carbon-emitting suppliers while satisfying growing customer

demand for green products. It's a win-win, but it's not normal practice yet.

Customers

The core thing companies are doing—and must continue to do—is helping customers reduce carbon emissions by developing and offering products that produce fewer emissions throughout their life cycles. We're seeing great innovation, and customer buy-in, for lower-footprint products in the biggest carbon-emitting sectors: electric vehicles in transportation; efficient heating, cooling, and lighting in buildings; and tasty alternative proteins in food and agriculture.

Manufacturers and retailers are also working to increase the use of recycled materials and reduce the amount of material used in packaging—all the way to zero in some cases. A group of British retailers, for example, has teamed up to change how some products leave the store. Consumers can fill their own bags and jars from bins of dry goods (grains, beans, nuts, and so on), laundry detergent, and shampoo. Some commercial products are trying to go even further: After making each tile of its prototype carbon-negative flooring, Interface explains, "there is less carbon dioxide in the atmosphere than if it had not been manufactured in the first place."

But businesses need to make products like these mainstream and then go beyond the direct impacts of their products on customers to drive deeper change. Here are three possible ways forward:

Help customers use less and mobilize. The two most aggressive actions companies can take with consumers are encouraging them to reduce consumption and engaging them in climate activism. Zurich-based Freitag, which makes bags from recycled materials, lets customers create a new look by switching bags with other customers. And Patagonia (always a radical company) is teaching its customers how to repair its clothes so that they don't need to buy new items. These companies may risk selling less, but they're building trusted brands with a loyal following. And discouraging consumption hasn't hurt Patagonia in the least: Sales have quadrupled over the past decade, reaching an estimated $1 billion. Going further, the company is using the trust it has built to mobilize consumers, through its Patagonia Action Works initiative, to engage with grassroots environmental groups in Europe and the United States.

Rising temperatures, rising risks: nature's collapse

If the global temperature were to increase by ...

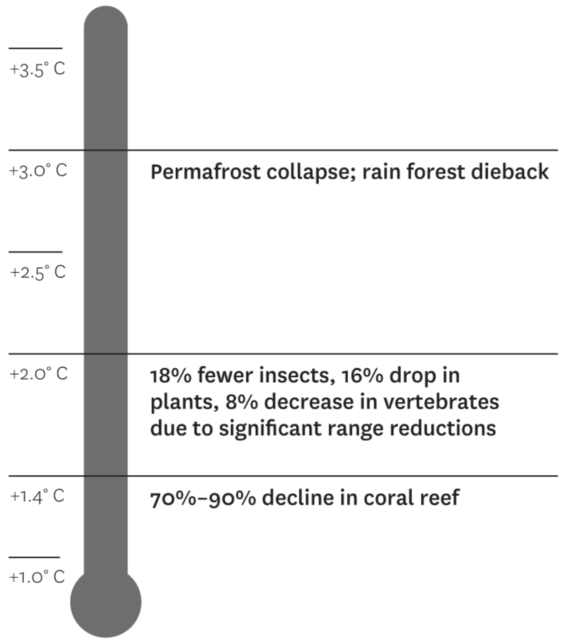

Source: World Resources Institute.

Use communications to educate and inspire consumers. Companies can make more-effective use of two channels in driving climate discussions: packaging and advertising. How? The Swedish oat drink brand Oatly, for example, reports product carbon emissions on its packages and points consumers to information on the climate benefits of eating plant-based products. Ben & Jerry's used the packaging and launch of an ice cream flavor, Save Our Swirled, to raise awareness about the Paris Climate Accords in 2015. IKEA surveyed more than 14,000 customers in 14 countries to understand their attitudes and how best to motivate climate action through advertising; the resulting framework is designed to guide its communications. In the fall of 2019 the household products company Seventh Generation donated advertising

airtime on the *Today* show to help promote the Youth Climate Movement.

A new collaborative initiative seeks to make promotional activities like these the norm. Launched recently by Sustainable Brands (on whose advisory board I sit)—along with some big names such as PepsiCo, Nestlé Waters, P&G, SC Johnson, and Visa—the Brands for Good program commits participants to encourage sustainable living through their marketing and communications and, even more ambitious, to transform the field of marketing to support that goal.

Choose business customers wisely. The efforts described above focus on traditional consumers. But companies need to direct equal attention to their business customers. As with suppliers, they must stop enabling customers that are either not addressing climate change or, more to the point, part of the high-carbon economy. Banks, venture capital and private equity funds, consulting companies, legal firms, and other service providers should ask tough questions about whom they're supporting. Helping companies be "better" at extracting or burning carbon-based fuels is actively moving the world in the wrong direction, and it dwarfs any carbon reduction a service business pursues in its own operations.

In the investment world, a movement to divest from fossil fuels is taking off, spearheaded by a group of investors with $11 trillion in assets. Norway's $1 trillion sovereign wealth fund is likewise dumping investments in many oil and gas companies.

Other service companies, such as consulting giants and law firms, that still work with carbon-intensive industries should be helping them make the permanent pivot necessary to survive. That means helping fossil fuel companies sunset their core business over the next few decades and completely shift their portfolios and business models toward clean options. Tech companies have to do some hard thinking as well. One of the reasons Amazon's employees rebelled was the company's announcement that its cloud business would help oil and gas companies accelerate exploration. Stakeholders will continue to ask probing questions about what companies stand for and whom they support—and companies will have to have an answer.

Employees

In the battle for talent, especially for Millennials and Gen Z, companies must prove that they are good citizens. Surveys consistently show that people under 40 want to work for employers that share their values. As Unilever's sustainable living plan gained steam in the mid-2010s, the company became the most sought-after employer in its sector. Top executives I've worked with at Unilever cite its sustainability leadership as key in attracting and retaining talent. The benefit flows both ways: Companies need their employees' commitment and buy-in to achieve their sustainability goals.

To reinforce this relationship, companies must build sustainability and climate action into their regular incentive

structures and systems—that is, pay everyone from the C-suite on down to cut carbon. They are secretive about the exact percentages, but the most committed companies I've seen tie at least a quarter of bonuses to sustainability key performance indicators (KPIs). It's time to increase that.

Can companies go even further and proactively support their employees' values by helping them drive change in the world around them? Some organizations already do. During the 2018 U.S. election, more than 100 of them, including Walmart, Levi Strauss, The Gap, Southwest Airlines, Kaiser Permanente, and Lyft, joined the Time to Vote initiative, giving employees time off to be good citizens. Some even encourage direct climate activism. Having identified the "climate emergency" as a top employee concern, the $1 billion cosmetics retailer Lush closed 200 shops in the U.S. to allow employees to join global climate marches last September. A Lush representative told me that during Canadian marches the company also shuttered 50 shops and offices for 20 manufacturing and support teams.

Atlassian, the fast-growing Australian enterprise software company with a $30 billion market cap, also encourages employees to become climate activists. As the company's cofounder Mike Cannon-Brookes wrote in his blunt blog "Don't @#$% the planet," Atlassian gives employees a week each year to volunteer for charity, and they can now use the time to join marches and strikes. He wants them to "go further and volunteer their time to other not-for-profit groups with a focus on climate."

Employees want to work for a company that stands for something. But they increasingly also want the freedom to express what *they* stand for. So ask them what they care about—especially younger and newer employees—and help them live their values.

3. Rethinking the Business

Flexing political muscle and reconceiving stakeholder relationships must happen quickly. But it is also time to think big, to look for new possibilities, and to question core assumptions about consumption and growth in the economy—that is, to go far beyond simply slashing energy use and buying renewables. Today the possibilities are broad, with everything from reducing food waste to developing circular business models falling under the umbrella of "climate strategy." Now is the right time to think critically and creatively about how *all* products and services in every sector are created and used and to squeeze carbon out of every step in the value chain. Some of this is tactical—for example, working with suppliers or customers to reduce their emissions, as discussed. But at the strategic level it can mean rethinking the company's investments and business models entirely. Here are some ways to do just that, focused on two key areas.

Risk and investments

Companies deploy capital and make investment decisions in multiple ways. With some important changes in how they think about financing and investment, much more capital could flow to low-carbon activities.

Rising temperatures, rising risks: heat waves

If the global temperature were to increase by …

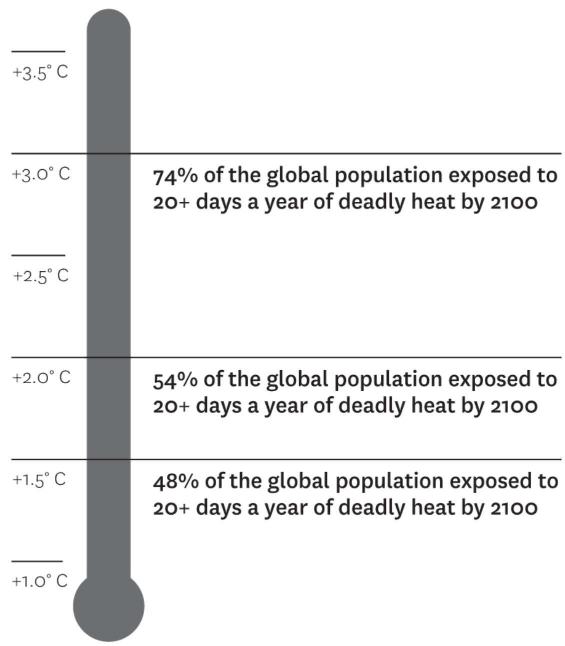

Source: World Resources Institute.

Note: According to research published in *Nature Climate Change*, "deadly heat" is the threshold beyond which air temperatures, humidity, and other factors can be lethal.

Consider the idea of return on investment. In most companies, to get internal funding, a project must achieve a predetermined rate of return (or hurdle rate) that will pay off relatively quickly. This approach to ROI is flawed. It generally measures the "R" in straight cash, without allowing for more-strategic or intangible

value. It's also agnostic as to whether the investment moves the company down a more sustainable path. We need to use this tool differently to shift to low-carbon investment choices.

Smart tweaks to two internal processes—capital expenditures and hurdle rates—can do a lot of good. J.M. Huber, a family-owned business that manufactures nature-based ingredients for the food and personal care industries along with components in home building, developed a more holistic approach to optimizing capital deployment. The chief sustainability officer and the CFO worked together to shift the capex process to factor in intangible benefits such as community engagement, customer perceptions, employee attraction and retention, and business resiliency (for example, solar array projects that insulate the business from fossil fuel energy price shocks).

Rising temperatures, rising risks: water uncertainty

If the global temperature were to increase by ...

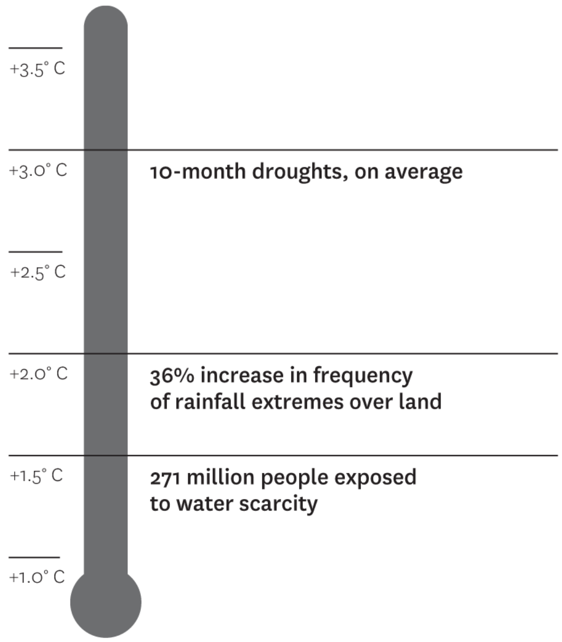

Source: World Resources Institute.

Note: According to the NOAA, "extreme rainfall" can be loosely defined as a month's worth of rain for a given region falling in a single day.

Companies should set their hurdle rates more strategically and allow some investments more leeway, with a strong bias toward funding carbon-reducing projects. If, for example, constructing an energy-efficient building—one that will save money and carbon over its lifetime—costs more up front or requires more than a few years to pay off, isn't it still a smart investment on a 40-year asset?

Another wise investment shift involves levying an internal carbon price on companies' own operations to encourage emissions reduction. More than 1,400 organizations now use internal pricing in some way, but the norm is to use "shadow" prices with no money changing hands. That approach isn't strong enough. Early leaders like Microsoft, Disney, and LVMH have

been collecting *real* money from divisions or functions related to their emissions. That "tax" revenue is reinvested in energy efficiency, renewables, or offset projects such as tree planting. All companies should use this strategy to help fund low-carbon projects and to prepare the business as government-imposed carbon taxes become more common.

A more recent strategy is to use financing tools such as green bonds, now a $200 billion market, in which the proceeds from bond purchases go to environmental and climate projects. The Italian energy group ENEL is trying something a bit different, issuing a bond tied to a KPI measuring the company's performance against the UN's Sustainable Development Goals. If ENEL misses its target of increasing renewable energy to 55% of its installed capacity, it will pay 25 basis points more to bondholders. Although the funds raised are not tied to a specific use, as they are with conventional green bonds, the instrument clearly supports emissions reduction.

Perhaps the biggest move a company can make is to rethink where to place its R&D bets. In a telling seismic shift, Daimler announced that it would no longer invest in research on internal combustion engines and would put billions toward electric vehicles instead. And the CEO of Nestlé, Mark Schneider, spoke recently about investing in plant-based proteins, which have a *much* smaller carbon footprint than conventionally produced meat, saying, "A Swiss franc we spend developing the burger is a burden to this quarter's profits. Next year or the year after, it will

come back to us if we do our job right." Seeing returns on a fast-growing new market within a year or two sounds like a good deal.

New business models

The level of carbon reduction that the Intergovernmental Panel on Climate Change says is required to head off catastrophic warming—cutting emissions in half by 2030 and to zero by 2050—is daunting. Everything discussed here will move us much more quickly, but some fundamental changes are needed in how we think about products, services, and consumption. Current business models and delivery methods can lock us into more material- and energy-intensive pathways. And some sectors, the most carbon-intensive, will need to exit core businesses.

Consider Philips Lighting, which launched a "light as a service" model, through which business customers pay Philips to install and manage their lighting rather than purchase a lighting system themselves. This flips Philips's traditional model on its head: Instead of trying to sell as many bulbs as possible, under this program, the company manages the provision of light as frugally as it can, using longer-lasting, more-efficient products that slash material and energy use. In a larger-scale transformation, the energy company Ørsted—formerly known as Danish Oil & Natural Gas—anticipated the decarbonization of the global economy and began pivoting from its core business a decade ago. It has since sold off most of its fossil fuel assets and has become the world's largest builder of offshore wind farms. And just a few

years ago, the idea that meat-based McDonald's and Burger King would both be selling plant-based "burgers" seemed far-fetched. But they, like Ørsted, may be thinking strategically about what the coming low-carbon economy means for their business.

The Next Level of Action

There's no doubt that companies are doing a lot on climate, including cutting emissions and setting aggressive carbon goals for operations, supply chains, and their innovation agendas. But it's not enough. The science is getting away from us, and we're losing the relatively stable planetary temperature range that allowed us to build our society over the past 10,000 years. Companies have many levers to pull to truly change business as usual, but most remain stuck in old thinking. Climate action is usually focused on incremental change. And even when they're setting a big goal like going to all-renewable energy, companies have waited until every project makes money quickly. Now they need to mobilize *all* corporate assets, hard *and* soft, to tackle this shared, unprecedented problem at the scale it requires.

Next-gen climate actions, as they become an expected part of business, will create significant long-term value. They will help companies build closer, lasting connections with key stakeholders; create clear and consistent regulatory environments that enable more sustainable practices that lower costs; and drive deeper, more-disruptive (or what I call *heretical*) innovation. Throw in the substantial intangible value—employee

attraction and loyalty, lowered risk in supply chain, resilience, license to operate, societal relevance, and preparation for a very different future—and you have a powerful business case.

But it's also well past time to recognize that aggressive climate action is necessary if humanity is to survive and thrive. Business and society won't succeed unless and until we do all we can to tackle climate change.

Your Company's Next Leader on Climate Is … the CFO

by Laura Palmeiro and Delphine Gibassier

If your chief financial officer is the last person you would think of to take charge on climate change, think again. Today, smart organizations are shifting their sustainability responsibilities toward the finance function.

There are several reasons for this change. First is the basic math, which falls largely within a CFO's purview. Mitigating and adapting to climate change will require close to $1 trillion in investments per year through 2030 for the economy as a whole, and is also expected to put at risk between $4.2 trillion and $43 trillion of tradable stock exchange assets by the end of the century, depending on the level of planetary warming. (The

latter number is for a world that has warmed by 6 degrees Celsius.)

Second, cutting greenhouse gas (GHG) emissions leads to cost savings. If you cut emissions, you cut energy, which is a massive organizational cost—something CFOs pay close attention to. Third, because investors are pushing to make climate-safe investments, they want climate risks to be integrated within corporate financial disclosures. Finally, the business opportunities for climate change solutions are blooming. According to Chartered Professional Accountants of Canada, "As creators, enablers, preservers and reporters of sustainable value, accountants can make their organizations' adaptation efforts more effective." Taken together, these shifts are leading finance teams to include what were formerly called "nonfinancials" in their daily jobs.

CFO leadership on climate change is starting to pay off. For example, Adnams, a British brewery, recently saw an increase in the base cost of beer because hot summers were affecting barley production. To solve the problem, the CFO was able to offset these higher costs by looking at energy and water savings. The CFO of Mars, Claus Aagaard, has talked about how the company's sustainability plan allowed it to capitalize on cost savings within two years.

Through our research, our corporate experience at Danone, and our work with the UN Global Compact, we have determined four key ways in which sustainability is being centralized in the

finance function—ways every corporate leader should be aware of.

Financial Tools Are Becoming More Green

Increasingly, we've seen finance teams greening more of their tools. What does this look like? Companies such as SSE or the Coca-Cola Hellenic Bottling Company, for example, have implemented "green CAPEX [capital expenditure]" systems. These structures, which involve small changes in investment decisions (like including an internal price on carbon emissions or loosening the payback period for investment decisions), have allowed climate change–friendly investments to take place on a larger scale.

Even more significant, Microsoft now has an internal carbon market codesigned by the finance and sustainability teams. Thanks to a carbon fee paid by subsidiaries based on the level of their GHG emissions—incentivizing them to cut their emissions—Microsoft has a carbon fund that fuels climate change–related investments, allowing more significant and global investments to be made. On January 16, 2020, Microsoft made a historic announcement, backed by its CFO, to become carbon negative by 2030 and remove their historical carbon emissions by 2050.

In fact, more than 600 organizations say they now use carbon pricing, for a number of different reasons, among them to inform procurement and R&D decisions, help suppliers transition to a low-carbon world, pay bonuses, or help with long-term

investments. In another change, Danone has started rewarding strong group performance by connecting incentives to climate change performance based on annual CDP scores.

Finally, following the integration of climate change within management control systems, corporations have started to measure GHG emissions like they measure their financials. Oracle has used what it calls “environmental accounting and reporting” to capture and transform GHG emissions from the company’s portfolio of 600 buildings across more than 70 countries. This has led to significant cost savings, because accurate data is being collected quickly. Even the small French company Saveurs et Vie, which produces food baskets for the elderly, has asked its enterprise resource planning system provider to allow it to automate carbon footprinting.

Finance Teams, Collaborations, and Roles Are Evolving

Changes in finance and accounting departments are increasingly visible within not only the tools but also the teams. Ørsted, a wind-power company based in Denmark, has a full-time environmental, social, and governance (ESG) accounting team made up of four employees. The UK-based energy provider SSE has a full-time sustainability accountant in-house. Since 2013, Unilever has had a finance director for sustainability, who is in charge of developing an understanding of sustainability in

finance, integrating sustainability into finance reporting, and developing best practices.

These company-specific examples are giving way to larger collaborations, too. The CFO Leadership Network, created in 2010 by Accounting for Sustainability in the UK, recently developed two Canadian and U.S. charters.

Some are rethinking the traditional CFO role altogether. In 2018, the Institute of Management Accountants published the first study on the emergence of sustainability CFOs (coauthored by one of us, Delphine), demonstrating the need for specific hybridized competencies between finance and sustainability to answer today's challenges. This research uncovered new competencies these leaders need to have, including developing natural capital profit and loss accounts, identifying the cost of key externalities, and understanding the value created through intangibles. Going further, Mervyn King (who is credited with the birth of "integrated reporting" in South Africa) developed the concept of a chief value officer in a 2016 book. And in North America, Manulife brought on a sustainability accounting director as a new kind of role.

Rules and Regulations Are Changing Rapidly

Your CFO will also need to adapt to shifting financial accounting rules that address climate change–related risks and opportunities. The biggest changes stem from December 2015, when the Financial Stability Board, an international body that

monitors and makes recommendations about the global financial system, established the Task Force on Climate-related Financial Disclosures (TCFD) “to develop a set of voluntary, consistent disclosure recommendations for use by companies in providing information to investors, lenders and insurance underwriters about their climate-related financial risks.” The new TCFD recommendations were released in June 2017 and included the suggestion that climate-related financial disclosures be made within mainstream annual financial filings and under governance processes similar to those for public disclosures.

What does this mean in practice? For one, all disclosures, including climate-related risks, climate metrics, and targets, should be reviewed by a company’s CFO, audit committee, or both. Companies also should face the future risks of their business models through scenario analysis.

In November 2019, the International Accounting Standards Board (IASB), whose mission is to develop accounting standards for financial markets around the world, published the report “IFRS Standards and Climate-Related Disclosures,” which recommended that companies address material environmental and societal issues and, more specifically, issues driven by investor pressure to disclose climate-related risks. (This was especially significant because the IASB usually does not mention climate change in accounting standards or briefings.) We expect recommendations like those from the TCFD and the IASB to continue.

The Financial Markets Increasingly Require a Focus on Climate

The financial markets are driving CFOs to look seriously at climate change. For example, the investor initiative Climate Action 100+, representing more than 370 investors with over $35 trillion in assets collectively, is urging 100 systemically important emitters to curb emissions, improve governance, and strengthen climate-related financial disclosures. Other initiatives, such as the climate benchmarks published by the European Union or the UN's Net Zero Asset Owner Alliance, are shifting the investment world into climate-ready financing. And in his annual letter to CEOs, BlackRock's Larry Fink emphasized that "the evidence on climate risk is compelling investors to reassess core assumptions about modern finance." Ultimately, Fink concluded that "climate risk is investment risk" and is alerting clients that BlackRock is centering its investment approach around sustainability.

Another reason for CFOs to take climate seriously comes from investors' appetite for green bonds—bonds that enable capital raising and investment for new and existing projects with environmental benefits. In 2019, new issuances on the green bond market reached around $250 billion overall, channeling more and more investments toward fighting climate change. Within this market, certified climate bonds, which are verified according to the type of physical asset or infrastructure they fund, allow companies to precisely align themselves with the 2015 Paris Agreement because they are consistent with its

warming limit of 2 degrees Celsius. In addition to enabling the financing of environmental projects, these instruments may even represent an advantage in terms of cost of capital, since external financing can, in some cases, become indexed on ESG performance.

When Peter Bakker from the World Business Council for Sustainable Development said in 2012 that "accountants would save the planet," he was not far from the truth. Today, accountants are increasingly prioritizing climate change inside their organizations and beyond. Your CFO should be the next leader to follow.

A Better Way to Talk about the Climate Crisis

by Gretchen Gavett

Many of us care about the climate, but it can be challenging to talk about. It's easy to get bogged down in stats and statistics, for one. And it can be nerve-racking to approach someone if you don't already know what their beliefs on the topic are. Sometimes, it's easier to just keep our mouths shut.

Given the urgency of the climate crisis, however, many of us feel that silence is no longer an option. And Dr. Katharine

Hayhoe, a climate scientist at Texas Tech University, is the person to talk to about how to talk about climate change. Hayhoe, whose 2018 TEDWomen talk on the subject has been viewed almost 2 million times, talks to everyone about the topic: Uber drivers, church ladies, Rotary Club members, business leaders, managers, elected officials, and more. People may have different backgrounds and views, but she's found a strategy that works: focusing on the heart—that is, what we collectively value—as opposed to the head.

So no matter your conversational goal, whether it's encouraging your company to act on climate issues or getting your employees to understand how the decisions they make affect your company's climate goals, this edited interview with Dr. Hayhoe is a great place to start.

What should any leader take into consideration when talking to people—employees, clients, suppliers, etc.—about climate change?

Ultimately, whether you're training a new employee, reviewing best practices with a supplier, or just having a conversation about climate change with a client, follow this rule of thumb: Don't start with fear, judgment, condemnation, or guilt. And don't start with just overwhelming people with facts and figures. *Do* start by connecting the dots to what is already important to both of us, and then offer positive, beneficial, and practical solutions that we can engage in.

Why have you found that this method works best? And how does it lead people toward understanding the urgency of climate change and taking action?

Often we believe that to care about climate change we have to be a certain type of person: an environmentalist, someone who bikes to work, or is a vegan. And if we're not any of those things, then we think, "Why should climate change matter to me?" But the reality is that if we are a human living on planet Earth, then climate change already matters to every single one of us; we just haven't realized it yet. Why? Because climate change affects the economy, the availability of natural resources, prices, jobs, international competition, and more. Failing to account for climate change in future long-range planning could lose us a competitive edge even in a best-case scenario, and potentially mean the end of a product line or an entire business in the worst case. By connecting climate impacts to what we already care about, we can recognize the importance and urgency of taking action.

So if I'm a leader, what are some specific ways in which I can communicate with my employees that sustainability is a key part of their jobs?

I would start early. During their initial training, I would explain very clearly how our products, our production, and our waste contributes to the problem of climate change. If our production is very energy intensive or produces a lot of organic waste, for example, that means we may be generating massive amounts of greenhouse gases. If our goods are transported over long distances, that also requires fossil fuels that produce heat-trapping gases. And aside from the issue of climate change, if we produce a lot of nonrecyclable waste that just piles up in landfills or the ocean, how much are we contributing to the pollution problem as well?

But I would also be sure to pair this information hand in hand with what we're doing to fix the problems from our end and how it's paying off. Give people analogies so it's really clear, so they can see it. I love giving examples of how many *X* worth of *Y* we've reduced; for example, something like "Through increasing the energy efficiency of our facilities, we have taken the equivalent of 500 cars off the road. Isn't that incredible? That's what we've been doing through our efforts." Or, "We have reduced our waste by 50%. That's the equivalent of *X* garbage trucks of waste per year." Or, "We are now powered by 38 wind turbines; that's *X* trainloads of coal we don't need to use anymore."

Finally—and this is the most important part!—I'd engage the employees themselves in the solutions. As humans, we want to be part of a solution. We want to make a difference. That is part of what gives us hope and what gives us energy, the idea that we're actually doing something good for the world.

So, for example, I might say, "We're aiming for an even better milestone. I want your ideas to help us get to this new milestone, too." That's even more incentivizing, when you feel like a company encourages you and supports you and wants you to be part of their plan.

Does this advice extend to people who might not believe that climate change is that severe—or that it exists at all? What might this kind of conversation look like in a professional setting?

Only around 10% of the population is dismissive [of climate change], but they are a very loud 10%. Glance at the comment section of any online article on climate change, check out the responses to my tweets, or search for global warming videos on YouTube—they're everywhere. They're even at our Thanksgiving dinner, because just about every one of us has at least one person who is dismissive in the family. I do, too!

A person who is dismissive is someone who has built their identity on rejecting the reality of a changing climate because they believe the solutions represent a direct and immediate threat to all they hold dear. And in pursuit of that goal, they will reject anything: hundreds of scientific studies, thousands of experts, even the evidence of their own eyes. So, no, there is no

point talking to a dismissive about climate science or impacts, unless you enjoy banging your head against the wall.

But it *can* be possible to have a constructive conversation with a dismissive—and I've had these!—by focusing solely on solutions that they don't see as a threat because they carry positive benefits and/or are good for their bottom line. And the fascinating thing is that once they are engaged in helping fix the problem, that very action can have the power to change a dismissive person's mind.

I want to end by asking about the importance of climate conversation over the next few years. I've heard anecdotally that companies are hearing more questions from younger job candidates or employees: "What are you doing? How are you addressing climate change as a company?" Does that resonate with you at all? Should companies be preparing for more conversations like these?

We see a very strong age gradient when it comes to levels of concern about climate change primarily among conservative populations, with younger people caring much more and being much more engaged than their elders. (Among more liberal populations, levels of concern are relatively high across all age groups.) At my own school, the number of students going to the president and asking, "What is our university doing?" has increased noticeably. I hear this anecdotally from colleagues all around the country, too. And when those students graduate, that's what they ask in their interviews, because they want to be part of the solution. Young people understand how urgent the

problem is, and they know that there's no time to waste. A lot of them don't want to do a job that is not helping to fix this massive problem that we have.

If companies want to be competitive, if they want to hire the best and the brightest, the ones who are most engaged, the ones who are most in tune, the ones who really put their heart and their soul and their passion into their work, then they have to start talking about climate change differently. Because this is increasingly becoming something that young professionals really care about.

Originally published in January 2020. Reprint BG2001

That Discomfort You're Feeling Is Grief

by Scott Berinato

SOME OF THE HBR EDIT STAFF met virtually the other day—a screen full of faces in a scene becoming more common everywhere. We talked about the content we're commissioning in this harrowing time of a pandemic and how we can help people. But we also talked about how we were feeling. One colleague mentioned that what she felt was grief. Heads nodded in all the panes.

If we can name it, perhaps we can manage it. We turned to David Kessler for ideas on how to do that. Kessler is the world's foremost expert on grief. He cowrote with Elisabeth Kübler-Ross *On Grief and Grieving: Finding the Meaning of Grief through the Five Stages of Loss*. His new book adds another stage to the process, *Finding Meaning: The Sixth Stage of Grief*. Kessler also has worked for a decade in a three-hospital system in Los Angeles. He served on their biohazards team. His volunteer work includes being an LAPD Specialist Reserve for traumatic events as well as having served on the Red Cross's disaster

services team. He is the founder of www.grief.com, which has over 5 million visits yearly from 167 countries.

Kessler shared his thoughts on why it's important to acknowledge the grief you may be feeling, how to manage it, and how he believes we will find meaning in it. The conversation is lightly edited for clarity.

People are feeling any number of things right now. Is it right to call some of what they're feeling grief?

Yes, and we're feeling a number of different griefs. We feel the world has changed, and it has. We know this is temporary, but it doesn't feel that way, and we realize things will be different. Just as going to the airport is forever different from how it was before 9/11, things will change and this is the point at which they changed. The loss of normalcy; the fear of economic toll; the loss of connection. This is hitting us and we're grieving. Collectively. We are not used to this kind of collective grief in the air.

You said we're feeling more than one kind of grief?

Yes, we're also feeling anticipatory grief. Anticipatory grief is that feeling we get about what the future holds when we're uncertain. Usually it centers on death. We feel it when someone gets a dire diagnosis or when we have the normal thought that we'll lose a parent someday. Anticipatory grief is also more broadly imagined futures. There is a storm coming. There's something bad out there. With a virus, this kind of grief

is so confusing for people. Our primitive mind knows something bad is happening, but you can't see it. This breaks our sense of safety. We're feeling that loss of safety. I don't think we've collectively lost our sense of general safety like this. Individually or as smaller groups, people have felt this. But all together, this is new. We are grieving on a micro and a macro level.

What can individuals do to manage all this grief?

Understanding the stages of grief is a start. But whenever I talk about the stages of grief, I have to remind people that the stages aren't linear and may not happen in this order. It's not a map, but it provides some scaffolding for this unknown world. There's denial, which we say a lot of early on: *This virus won't affect us.* There's anger: *You're making me stay home and taking away my activities.* There's bargaining: *Okay, if I social distance for two weeks, everything will be better, right?* There's sadness: *I don't know when this will end.* And finally there's acceptance. *This is happening; I have to figure out how to proceed.*

Idea in Brief

During the global pandemic, a palpable sense of collective grief has emerged. In an interview with HBR, grief expert David Kessler explains the classic five stages of grief and the practical steps we can take to manage these emotions throughout the crisis. These include balancing bad thoughts with good; focusing on the present; letting go of things you can't control; and stocking up on compassion. Kessler also talks about a sixth stage of grief: meaning. After acceptance, he

says, we will find meaning in the hard-to-fathom events, and we will be stronger for it.

Acceptance, as you might imagine, is where the power lies. We find control in acceptance. *I can wash my hands. I can keep a safe distance. I can learn how to work virtually.*

When we're feeling grief, there's that physical pain. And the racing mind. Are there techniques to deal with that to make it less intense?

Let's go back to anticipatory grief. Unhealthy anticipatory grief is really anxiety, and that's the feeling you're talking about. Our mind begins to show us images. My parents getting sick. We see the worst scenarios. That's our minds being protective. Our goal is not to ignore those images or to try to make them go away—your mind won't let you do that, and it can be painful to try and force it. The goal is to **find balance in the things you're thinking**. If you feel the worst image taking shape, make yourself think of the best image. We all get a little sick and the world continues. Not everyone I love dies. Maybe no one does because we're all taking the right steps. Neither scenario should be ignored, but neither should dominate either.

Anticipatory grief is the mind going to the future and imagining the worst. To calm yourself, you want to **come into the present**. This will be familiar advice to anyone who has meditated or practiced mindfulness, but people are always

surprised at how prosaic this can be. You can name five things in the room. There's a computer, a chair, a picture of the dog, an old rug, and a coffee mug. It's that simple. Breathe. Realize that in the present moment, nothing you've anticipated has happened. In this moment, you're okay. You have food. You are not sick. Use your senses and think about what they feel. The desk is hard. The blanket is soft. I can feel the breath coming into my nose. This really will work to dampen some of that pain.

You can also think about how to **let go of what you can't control.** What your neighbor is doing is out of your control. What is in your control is staying six feet away from them and washing your hands. Focus on that.

Finally, it's a good time to **stock up on compassion.** Everyone will have different levels of fear and grief, and it manifests in different ways. A coworker got very snippy with me the other day, and I thought, *That's not like this person; that's how they're dealing with this. I'm seeing their fear and anxiety.* So be patient. Think about who someone usually is and not who they seem to be in this moment.

One particularly troubling aspect of this pandemic is the open-endedness of it.

This is a temporary state. It helps to say it. I worked for 10 years in the hospital system. I've been trained for situations like this. I've also studied the 1918 flu pandemic. The

precautions we’re taking are the right ones. History tells us that. This is survivable. We will survive. This is a time to overprotect but not overreact.

And, I believe we will find meaning in it. I’ve been honored that Elisabeth Kübler-Ross’s family has given me permission to add a sixth stage to grief: meaning. I had talked to Elisabeth quite a bit about what came after acceptance. I did not want to stop at acceptance when I experienced some personal grief. I wanted meaning in those darkest hours. And I do believe we find light in those times. Even now people are realizing they can connect through technology. They are not as remote as they thought. They are realizing they can use their phones for long conversations. They’re appreciating walks. I believe we will continue to find meaning now and when this is over.

What do you say to someone who’s read all this and is still feeling overwhelmed with grief?

Keep trying. There is something powerful about naming this as grief. It helps us feel what’s inside of us. So many have told me in the past week, “I’m telling my coworkers I’m having a hard time,” or “I cried last night.” When you name it, you feel it and it moves through you. Emotions need motion. It’s important we acknowledge what we go through. One unfortunate by-product of the self-help movement is we’re the first generation to have feelings about our feelings. We tell ourselves things like, *I feel sad, but I shouldn’t feel that; other*

people have it worse. We can—we should—stop at the first feeling. *I feel sad. Let me go for five minutes to feel sad.* Your work is to feel your sadness and fear and anger whether or not someone else is feeling something. Fighting it doesn't help because your body is producing the feeling. If we allow the feelings to happen, they'll happen in an orderly way, and it empowers us. Then we're not victims.

In an orderly way?

Yes. Sometimes we try not to feel what we're feeling because we have this image of a "gang of feelings." If I feel sad and let that in, it'll never go away. The gang of bad feelings will overrun me. The truth is a feeling that moves through us. We feel it and it goes and then we go to the next feeling. There's no gang out to get us. It's absurd to think we shouldn't feel grief right now. Let yourself feel the grief and keep going.

Originally published in March 2020. Reprint H05HVE

About the Contributors

SCOTT BERINATO is a senior editor at *Harvard Business Review* and the author of *Good Charts Workbook: Tips, Tools, and Exercises for Making Better Data Visualizations* (Harvard Business Review Press, 2019) and *Good Charts: The HBR Guide to Making Smarter, More Persuasive Data Visualizations* (Harvard Business Review Press, 2016).

MARCUS BUCKINGHAM is the head of people and performance research at the ADP Research Institute. He is a coauthor of *Nine Lies about Work: A Freethinking Leader's Guide to the Real World* (Harvard Business Review Press, 2019).

PETER CAPPELLI is the George W. Taylor Professor of Management at the Wharton School and the director of its Center for Human Resources. His most recent book is *Will College Pay Off? A Guide to the Most Important Financial Decision You'll Ever Make* (PublicAffairs, 2015).

TIZIANA CASCIARO is a professor of organizational behavior at the University of Toronto's Rotman School of Management.

PAIGE COHEN is an associate editor at *Harvard Business Review*.

AMY C. EDMONDSON is the Novartis Professor of Leadership and Management at Harvard Business School and author of

The Fearless Organization (Wiley, 2019).

TIM FOUNTAINE is a partner in McKinsey's Sydney office and leads QuantumBlack, an advanced analytics firm owned by McKinsey, in Australia.

GRETCHEN GAVETT is a senior editor at *Harvard Business Review*.

DELPHINE GIBASSIER is an associate professor of accounting for sustainable development at Audencia Business School, with 18 years of experience in financial and nonfinancial accounting. She holds a PhD from HEC Paris.

ASHLEY GOODALL is the senior vice president of leadership and team intelligence at Cisco Systems. He is a coauthor of *Nine Lies about Work: A Freethinking Leader's Guide to the Real World* (Harvard Business Review Press, 2019).

ANDREI HAGIU is an associate professor of information systems at Boston University's Questrom School of Business. Twitter: @theplatformguy.

DANE E. HOLMES is the global head of human capital management at Goldman Sachs.

SUJIN JANG is an assistant professor of organizational behavior at INSEAD.

ANTHONY J. MAYO is the Thomas S. Murphy Senior Lecturer of Business Administration in the Organizational Behavior Unit at Harvard Business School. With Laura Morgan Roberts and Morehouse College president David A. Thomas, he edited *Race, Work, and Leadership: New Perspectives on the Black Experience* (Harvard Business Review Press, 2019), an exploration of the resources and actions needed for organizations to eliminate systemic racism.

BRIAN MCCARTHY is a partner in McKinsey's Atlanta office and coleads the knowledge development agenda for McKinsey Analytics.

LAURA PALMEIRO is the senior adviser to the United Nations Global Compact. She has extensive experience in finance, controlling, and sustainability at PwC and Danone. She holds an MBA from IAE Argentina.

JENNIFER PETRIGLIERI is an associate professor of organizational behavior at INSEAD and the author of *Couples That Work: How Dual-Career Couples Can Thrive in Love and Work* (Harvard Business Review Press, 2019). At INSEAD she directs the Management Acceleration Programme, the Women Leaders Programme, and the Gender Diversity Programme.

GARY P. PISANO is the Harry E. Figgie Jr. Professor of Business Administration and the senior associate dean of faculty development at Harvard Business School. He is the author of *Creative Construction: The DNA of Sustained Innovation* (PublicAffairs, 2019).

LAURA MORGAN ROBERTS is a professor of practice at the University of Virginia's Darden School of Business. With Anthony J. Mayo and Morehouse College president David A. Thomas, she edited *Race, Work, and Leadership: New Perspectives on the Black Experience* (Harvard Business Review Press, 2019), an exploration of the resources and actions needed for organizations to eliminate systemic racism.

TAMIM SALEH is a senior partner in McKinsey's London office and heads McKinsey Analytics in Europe.

KATINA B. SAWYER is an assistant professor of management at The George Washington University.

NICOLAJ SIGGELKOW is a professor of management and strategy at Wharton and a codirector of the Mack Institute for Innovation Management. He is a coauthor of *Connected Strategy: Building Continuous Customer Relationships for Competitive Advantage* (Harvard Business Review Press, 2019).

JANE STEVENSON, the global leader for CEO succession and vice chair of board and CEO services at Korn Ferry, is a coauthor of *Breaking Away: How Great Leaders Create Innovation That Drives Sustainable Growth—and Why Others Fail* (McGraw-Hill, 2011).

CHRISTIAN TERWIESCH is a professor of operations and innovation at Wharton and a codirector of the Mack Institute for Innovation Management. He is a coauthor of *Connected Strategy: Building Continuous Customer Relationships for Competitive Advantage* (Harvard Business Review Press, 2019).

MELISSA THOMAS-HUNT holds a faculty appointment at Stanford Graduate School of Business and has worked in professorial roles for more than 20 years at schools such as Vanderbilt, University of Virginia, Cornell University, and Washington University in St. Louis. Her research and teaching has focused on organizational behavior and the factors that unleash, leverage, and amplify the talents and contributions made by women and underrepresented individuals. In May 2019, Dr. Thomas-Hunt joined Airbnb as head of global diversity and belonging. She serves on the company's executive team and leads the strategy and execution of the global internal diversity, inclusion, equity, and belonging programs.

CHRISTIAN N. THOROUGHGOOD is an assistant professor of psychology in the Department of Psychological and Brain Sciences and the graduate programs in human resource development at Villanova University.

JENNICA R. WEBSTER is a codirector of the Institute for Women's Leadership and an associate professor of management at Marquette University.

ANDREW WINSTON is the founder of Winston Eco-Strategies and a global expert on business and sustainability. He is the author of *The Big Pivot: Radically Practical Strategies for a Hotter, Scarcer, and More Open World* (Harvard Business Review Press, 2014).

JULIAN WRIGHT is a professor of economics at the National University of Singapore.

Index

Engage with HBR content the way you want, on any device.

With HBR's new subscription plans, you can access world-renowned **case studies** from Harvard Business School and receive **four free eBooks**. Download and customize prebuilt **slide decks and graphics** from our **Visual Library**. With HBR's archive, top 50 best-selling articles, and five new articles every day, HBR is more than just a magazine.

Subscribe Today
hbr.org/success

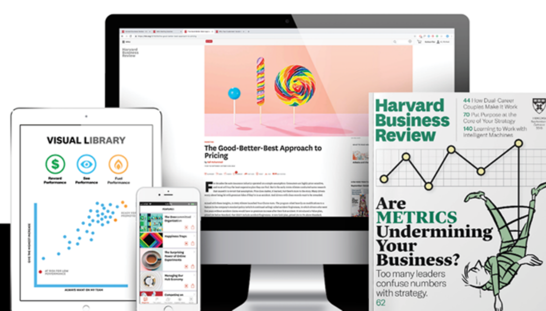

Harvard Business Review

The most important management ideas all in one place.

We hope you enjoyed this book from *Harvard Business Review*. Now you can get even more with HBR's 10 Must Reads Boxed Set. From books on leadership and strategy to managing yourself and others, this 6-book collection delivers articles on the most essential business topics to help you succeed.

HBR's 10 Must Reads Series

The definitive collection of ideas and best practices on our most sought-after topics from the best minds in business.

- Change Management
- Collaboration
- Communication
- Emotional Intelligence
- Innovation
- Leadership
- Making Smart Decisions
- Managing Across Cultures
- Managing People
- Managing Yourself
- Strategic Marketing
- Strategy
- Teams
- The Essentials

hbr.org/mustreads

Buy for your team, clients, or event.
Visit hbr.org/bulksales for quantity discount rates.

Harvard Business Review Press

HBR'S 10 MUST READS

BONUS ARTICLE

"Begin with Trust"

By Frances X. Frei and Anne Morriss

The definitive management ideas of the year from **Harvard Business Review.**

2022

HBR's 10 Must Reads series is the definitive collection of ideas and best practices for aspiring and experienced leaders alike. These books offer essential reading selected from the pages of *Harvard Business Review* on topics critical to the success of every manager.

Titles include:

HBR's 10 Must Reads 2015
HBR's 10 Must Reads 2016
HBR's 10 Must Reads 2017
HBR's 10 Must Reads 2018
HBR's 10 Must Reads 2019
HBR's 10 Must Reads 2020
HBR's 10 Must Reads 2021
HBR's 10 Must Reads for CEOs
HBR's 10 Must Reads for New Managers
HBR's 10 Must Reads on AI, Analytics, and the New Machine Age
HBR's 10 Must Reads on Boards
HBR's 10 Must Reads on Building a Great Culture
HBR's 10 Must Reads on Business Model Innovation
HBR's 10 Must Reads on Career Resilience

HBR's 10 Must Reads on Change Management (Volumes 1 and 2)

HBR's 10 Must Reads on Collaboration

HBR's 10 Must Reads on Communication (Volumes 1 and 2)

HBR's 10 Must Reads on Creativity

HBR's 10 Must Reads on Design Thinking

HBR's 10 Must Reads on Diversity

HBR's 10 Must Reads on Emotional Intelligence

HBR's 10 Must Reads on Entrepreneurship and Start-Ups

HBR's 10 Must Reads on Innovation

HBR's 10 Must Reads on Leadership (Volumes 1 and 2)

HBR's 10 Must Reads on Leadership for Healthcare

HBR's 10 Must Reads on Leadership Lessons from Sports

HBR's 10 Must Reads on Lifelong Learning

HBR's 10 Must Reads on Making Smart Decisions

HBR's 10 Must Reads on Managing Across Cultures

HBR's 10 Must Reads on Managing in a Downturn, Expanded Edition

HBR's 10 Must Reads on Managing People (Volumes 1 and 2)

HBR's 10 Must Reads on Managing Risk

HBR’s 10 Must Reads on Managing Yourself (Volumes 1 and 2)
HBR’s 10 Must Reads on Mental Toughness
HBR’s 10 Must Reads on Negotiation
HBR’s 10 Must Reads on Nonprofits and the Social Sectors
HBR’s 10 Must Reads on Organizational Resilience
HBR’s 10 Must Reads on Platforms and Ecosystems
HBR’s 10 Must Reads on Public Speaking and Presenting
HBR’s 10 Must Reads on Reinventing HR
HBR’s 10 Must Reads on Sales
HBR’s 10 Must Reads on Strategic Marketing
HBR’s 10 Must Reads on Strategy (Volumes 1 and 2)
HBR’s 10 Must Reads on Strategy for Healthcare
HBR’s 10 Must Reads on Teams
HBR’s 10 Must Reads on Women and Leadership
HBR’s 10 Must Reads: The Essentials

HBR'S 10 MUST READS

The definitive management ideas of the year from **Harvard Business Review.**

2022

HARVARD BUSINESS REVIEW PRESS

Boston, Massachusetts

Copyright

HBR Press Quantity Sales Discounts

Harvard Business Review Press titles are available at significant quantity discounts when purchased in bulk for client gifts, sales promotions, and premiums. Special editions, including books with corporate logos, customized covers, and letters from the company or CEO printed in the front matter, as well as excerpts of existing books, can also be created in large quantities for special needs.

For details and discount information for both print and ebook formats, contact booksales@harvardbusiness.org, tel. 800-988-0886, or www.hbr.org/bulksales.

Copyright 2021 Harvard Business School Publishing Corporation

All rights reserved

No part of this publication may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form, or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the prior

permission of the publisher. Requests for permission should be directed to permissions@harvardbusiness.org, or mailed to Permissions, Harvard Business School Publishing, 60 Harvard Way, Boston, Massachusetts 02163.

First eBook Edition: Oct 2021

ISBN: 978-1-64782-213-2
eISBN: 978-1-64782-214-9

Contents

Index

Editors' Note

As our editorial team read through the past year's issues of *Harvard Business Review* to select the articles for this volume, a feeling of hope emerged among us. After many trying months of social distancing, adjusting to remote work, political unrest, a worldwide recession, and illuminating social justice movements, vaccine rollouts were ramping up globally, and a brighter future felt within reach. When we chose these pieces, we tried not only to look back at the year that was but also to imagine what leaders would need at the dawn of 2022 to thrive in the new normal.

More often than ever before, articles HBR published were a direct response to current events. We covered the issue of anxiety and work for the first time because the alarming increase of people experiencing anxiety became impossible for leaders to ignore. The crises pushed many people to rethink their careers, what they expect from their employers, and how they'd like to work—making it a prime time for brushing up on job negotiation skills. These articles reminded us that life

would keep going on in spite of the chaos. We saw that in explorations of compassionate leadership, the best ways to grow an innovative culture, and the continuing rise of machine learning. Other pieces looked at trends already underway that had accelerated over the past year. Some organizations were working remotely before 2020, but the pandemic created a need to swiftly get up to speed with the right technology and the best team practices for working from anywhere. And although inequality and bias have been unresolved workplace issues for centuries, recent events have forced us to reexamine what it means to have a truly inclusive, diverse workplace—and to reflect on what hasn't worked in the past.

The challenges of 2020 revealed where we should start making changes for a better year ahead. The old ways aren't coming back, and in many cases we wouldn't want them to. As you read through this collection, we hope these articles inspire hope and help you feel prepared for what's next.

We start off with **"Begin with Trust,"** because as Frances Frei and Anne Morriss say, "Trust is the basis for almost everything we do." It's the reason we're

willing to exchange our hard-earned paychecks for goods and services, to pledge our lives to another person in marriage, and to cast a ballot for someone who will represent our interests. Trust is also essential for leaders who strive to empower other people as a result of their presence and to ensure that the impact of their leadership continues beyond their tenure. The more trust you build, the easier it is to practice that kind of leadership. How can leaders start? By being authentic, exercising sound judgment, and showing empathy. Frei and Morriss explain how to assess your strengths and weaknesses when it comes to trust and offer advice for building the three components of trust.

In **"Cultural Innovation,"** Douglas Holt argues that companies struggle with innovation because they put all their chips on one innovation paradigm—what he calls *better mousetraps*. Fortunately, there is a better way to innovate. In consumer markets, innovation often proceeds according to a logic Holt calls *cultural innovation*. With better-mousetraps innovation, companies focus on outdoing competitors on existing notions of value. With cultural innovation, you change the understanding of "valuable." Using the stories of

the Ford Explorer's reinvention of the family car and how Blue Buffalo turned the culture of dog food on its head, Holt reveals the strategic principles that allow companies to pursue cultural innovation.

"Co-opetition"—cooperating with a competitor to achieve a common goal or get ahead—has been gaining traction for three decades. Yet many companies are still uncomfortable with the concept and miss the promising opportunities it presents. In **"The Rules of Co-opetition,"** Adam Brandenburger and Barry Nalebuff offer a framework for deciding whether to team up with a rival and how to manage the relationship, drawing on examples from DHL, UPS, Google, and Yahoo. Co-opetition requires teams to think both competitively and cooperatively—at the same time. Firms that learn to do so can gain an important edge.

Cooperation is also key to successful negotiation. In **"Negotiating Your Next Job,"** Hannah Riley Bowles and Bobbi Thomason encourage readers seeking to advance their careers to think strategically about not just what they want but how to get it. The authors draw on their work coaching executives and their cross-cultural

research to propose four steps that can prepare you to negotiate: Think broadly about your long-term career goals, be mindful of what type of opportunity you're asking for, arm yourself with the necessary information, and connect with people who can be helpful in making your case. Be specific and realistic, they argue, and you're more likely to achieve success.

Although many workers dealt with anxiety before 2020, for some the uncertainty and stress of the pandemic brought it on for the first time. In "**Leading Through Anxiety**," Morra Aarons-Mele starts a productive, much-needed conversation about anxiety in the workplace and how we can reframe its negative reputation. She explains that although anxiety is uncomfortable, it isn't always counterproductive. Because managing anxiety can make us more comfortable with uncomfortable feelings, it can prompt us to react quickly to threats. And when channeled thoughtfully, it can make us better leaders in a crisis. But unchecked, it zaps energy and clouds decisions. In this article Aarons-Mele offers advice on how to inspire and encourage your team even when you're struggling yourself.

Products and services that rely on machine learning don't always lead to ethical or accurate choices. Sometimes they cause investment losses, for instance, or biased hiring, or car accidents. And as machine-learning-based AI offerings proliferate across markets, the companies creating them face major new risks. Executives need to understand and mitigate the technology's potential downside. In "**When Machine Learning Goes Off the Rails**," Boris Babic, I. Glenn Cohen, Theodoros Evgeniou, and Sara Gerke provide a guide to managing the risks.

Twenty-five years ago Robin J. Ely and David A. Thomas made the groundbreaking argument that to fully benefit from increased racial and gender diversity, organizations must adopt a learning orientation and be willing to change the corporate culture and power structure. In the time since, organizations have largely failed to do so and are no closer to reaping diversity's benefits. Instead, business leaders misconstrue or ignore what abundant research has made clear: Increasing the numbers of traditionally underrepresented people in your workforce does not automatically produce good results. Now, in "**Getting**

Serious About Diversity," the authors reevaluate and update their argument and highlight the key actions for leaders looking to make real change.

As intractable as it may seem, racism in the workplace can be effectively addressed with the right information, incentives, and investment. Robert Livingston argues that because organizations are small, autonomous entities that afford leaders a high level of control over norms and policies, they are ideal sites for doing so. In **"How to Promote Racial Equity in the Workplace,"** Livingston walks readers through the five stages of a process for making profound and sustainable progress toward this goal in your organization.

The pandemic has hastened a rise in remote working for knowledge-based organizations. This has notable benefits: Companies can save on real estate costs, hire and utilize talent globally, mitigate immigration issues, and experience productivity gains, while workers can enjoy geographic flexibility. At the same time, concerns include how to communicate across time zones, share knowledge that isn't yet codified, socialize virtually and prevent professional isolation, protect client data,

and avoid slacking. In “**Our Work-from-Anywhere Future**,” Prithwiraj (Raj) Choudhury highlights best practices that can help leaders decide whether remote working is right for their organizations.

“**A More Sustainable Supply Chain**” tackles the disconnect between the commitment of multinational corporations (MNCs) to sustainability and the reality of their suppliers. MNCs are pledging to procure the materials and services they need from companies committed to fair labor practices and environmental protections. But their suppliers—and their suppliers' suppliers—often violate sustainability standards, exposing MNCs to serious financial and social risks. To explore this problem—and identify solutions—Verónica H. Villena and Dennis A. Gioia studied the supply networks of three MNCs deemed to be sustainability leaders.

To close this collection, we take a fascinating look at a widely admired company whose internal workings have often been unknowable to outsiders. In “**How Apple Is Organized for Innovation**,” Joel M. Podolny and Morten T. Hansen discuss the innovation benefits and leadership challenges of Apple's distinctive and

ever-evolving organizational model in the belief that it may be useful for other companies competing in rapidly changing environments. When Steve Jobs returned to Apple, in 1997, it had a conventional structure for a company of its size and scope. It was divided into business units, each with its own P&L responsibilities. Believing that conventional management had stifled innovation, Jobs laid off the general managers of all the business units (in a single day), put the entire company under one P&L, and combined the disparate functional departments of the business units into one functional organization. Although such a structure is common for small entrepreneurial firms, Apple—remarkably—retains it today, even though the company is nearly 40 times as large in terms of revenue and far more complex than it was in 1997.

Keeping up with business trends and synthesizing the best ideas is important—and time-consuming—work for today's leaders. With this volume, we've done the heavy lifting for you. Despite the challenges of the past year, you can learn much from these articles as you lead your business forward. We hope they prepare

you for a better year ahead and set you on the right track for a prosperous future.

—The Editors

Begin with Trust

by Frances Frei and Anne Morriss

ON A SPRING AFTERNOON IN 2017, Travis Kalanick, then the CEO of Uber, walked into a conference room at the company's Bay Area headquarters. One of us, Frances, was waiting for him. Meghan Joyce, the company's general manager for the United States and Canada, had reached out to us, hoping that we could guide the company as it sought to heal from a series of deep, self-inflicted wounds. We had a track record of helping organizations, many of them founder-led, tackle messy leadership and culture challenges.

We were skeptical about Uber. Everything we'd read about the company suggested it had little hope of redemption. At the time, the company was an astonishingly disruptive and successful start-up, but its success seemed to have come at the price of basic decency.

In early 2017, for example, when taxi drivers went on strike in New York City to protest President Trump's travel ban, Uber appeared to have used tactics to profit from the situation—a move that prompted widespread outrage and a #deleteUber campaign. A month later, not long before the meeting, an Uber engineer named Susan Fowler had blogged courageously about her experiences of harassment and discrimination at the company, which caused more outrage. Footage of Kalanick had then emerged, in a video that went viral, of his interaction with an Uber driver, where he appeared dismissive of the pain of earning a living in a post-Uber world. Additional charges leveled at the company in this period reinforced Uber's reputation as a cold-blooded operator that would do almost anything to win.

Despite our skepticism, Frances had gone to California to hear Kalanick out. (Anne was building her own company at the time, so she took a back seat on the project.) As Frances waited for him to make his entrance, she braced herself for the smug CEO she'd read about. But that wasn't who walked in. Kalanick arrived humbled and introspective. He had thought a lot about how the cultural values he'd instilled in the company—the very values that had fueled

Uber's success—had also been misused and distorted on his watch. He expressed deep respect for what his team had achieved but also acknowledged that he'd put some people in leadership roles without giving them the training or mentorship to be effective. Whatever mistakes Kalanick had made up to that point, he revealed a sincere desire to do the right thing as a leader.

We regrouped back in Cambridge, Massachusetts, and debated whether to take on the project. There were lots of reasons to stay far away from it. The work would be hard and its outcome uncertain, to say nothing of the brutal commute. Uber's workforce was frustrated, and the brand was becoming toxic. But we realized that if we could help get Uber back on the right path, then we could offer a road map to countless others trying to restore humanity to organizations that had lost their way. So we signed on.

After making that decision, we knew exactly where to start. With trust.

Empowerment Leadership

We think of trust as precious, and yet it's the basis for almost everything we do as civilized people. Trust is the reason we're willing to exchange our hard-earned

paychecks for goods and services, pledge our lives to another person in marriage, cast a ballot for someone who will represent our interests. We rely on laws and contracts as safety nets, but even they are ultimately built on trust in the institutions that enforce them. We don't know that justice will be served if something goes wrong, but we have enough faith in the system that we're willing to make high-stakes deals with relative strangers.

Trust is also one of the most essential forms of capital a leader has. Building trust, however, often requires thinking about leadership from a new perspective. The traditional leadership narrative is all about you: your vision and strategy; your ability to make the tough calls and rally the troops; your talents, your charisma, your heroic moments of courage and instinct. But leadership really isn't about you. It's about empowering other people as a result of your presence, and about making sure that the impact of your leadership continues into your absence.

Idea in Brief

The Starting Point

The traditional leadership narrative is all about you: your talents, charisma, and moments of courage and instinct. But real leadership is about your people and creating the conditions

for them to fully realize their own capacity and power. To do this, you have to develop stores of trust.

The Challenge

How do leaders build trust? By focusing on its core drivers: authenticity, logic, and empathy. People tend to trust you when they think they're interacting with the real you, when they have faith in your judgment and competence, and when they believe you care about them.

The Way Forward

When leaders have trouble with trust, it's usually because they're weak on one of those three drivers. To develop or restore trust, identify which driver you're "wobbly" on, and then work on strengthening it.

That's the fundamental principle we've learned in the course of dedicating our careers to making leaders and organizations better. Your job as a leader is to create the conditions for your people to fully realize their own capacity and power. And that's true not only when you're in the trenches with them but also when you're not around and even—this is the cleanest test—when you've permanently moved on from the team. We call it empowerment leadership. The more trust you build, the more possible it is to practice this kind of leadership.

The Core Drivers of Trust

So how do you build up stores of this foundational leadership capital? In our experience, trust has three core drivers: authenticity, logic, and empathy. People tend to trust you when they believe they are interacting with the real you (authenticity), when they have faith in your judgment and competence (logic), and when they feel that you care about them (empathy). When trust is lost, it can almost always be traced back to a breakdown in one of these three drivers.

People don't always realize how the information (or more often, the misinformation) that they're broadcasting may undermine their own trustworthiness. What's worse, stress tends to amplify the problem, causing people to double down on behaviors that make others skeptical. For example, they might unconsciously mask their true selves in a job interview, even though that's precisely the type of less-than-fully-authentic behavior that reduces their chance of being hired.

The good news is that most of us generate a stable pattern of trust signals, which means a small change in behavior can go a long way. In moments when trust is broken, or fails to get any real traction, it's usually the same

driver that has gone wobbly on us—authenticity, empathy, or logic. We call this driver your "trust wobble." In simple terms, it's the driver that's most likely to fail you.

Everybody, it turns out, has a trust wobble. To build trust as a leader, you first need to figure out what yours is.

Build It, and They Will Come

To identify your wobble, think of a recent moment when you were not trusted as much as you wanted to be. Maybe you lost an important sale or didn't get a stretch assignment. Maybe someone simply doubted your ability to execute. With that moment in mind, do something hard: Give the other person in your story the benefit of the doubt. Let's call that person your "skeptic." Assume that your skeptic's reservations were valid and that you were the one responsible for the breakdown in trust. This exercise only works if you own it.

If you had to choose from our three trust drivers, which would you say went wobbly on you in this situation? Did your skeptic feel you were misrepresenting some part of yourself or your story? If so, that's an authenticity problem. Did your skeptic feel you might be putting your own interests first? If so, that's an empathy problem. Did your

skeptic question the rigor of your analysis or your ability to execute on an ambitious plan? If so, that's a logic problem.

The trust triangle

Trust has three drivers: authenticity, logic, and empathy. When trust is lost, it can almost always be traced back to a breakdown in one of them. To build trust as a leader, you first need to figure out which driver you "wobble" on.

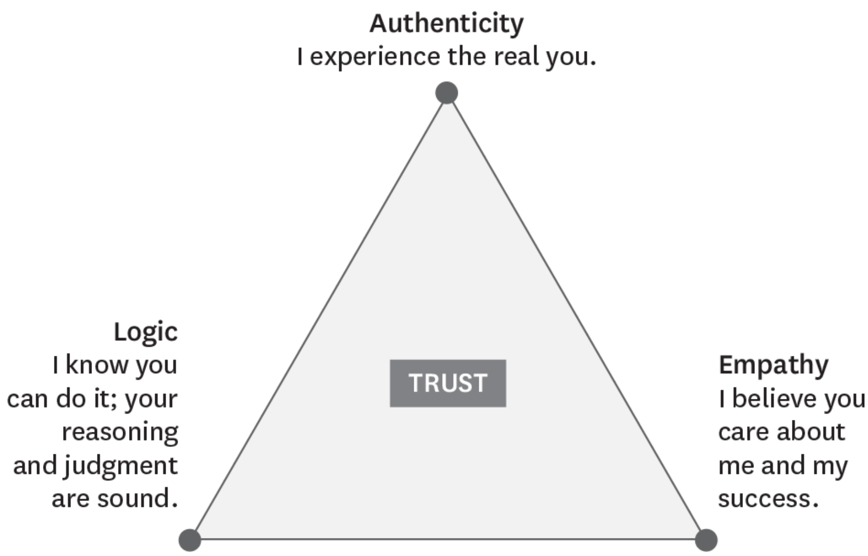

Now stand back and try to look at your pattern of wobbles across multiple incidents. Pick three or four interactions that stand out to you, for whatever reason, and do a quick trust diagnostic for each one. What does your typical wobble seem to be? Does the pattern change under stress or with different kinds of stakeholders? For example, do you wobble on one trait with your direct

reports but on a different one with people who have authority over you? That's not uncommon.

This exercise works best if you bring at least one person along for your diagnostic ride, ideally someone who knows you well. Sharing your analysis can be clarifying—even liberating—and will help you test and refine your hypothesis. In our experience, about 20% of self-assessments need a round of revision, so choose a partner who can keep you honest. Consider going back and testing your analysis directly by speaking openly about it with your skeptic. This conversation alone can be a powerful way to rebuild trust. When you take responsibility for a wobble, you reveal your humanity (authenticity) and analytic chops (logic) while communicating your commitment to the relationship (empathy).

Overcoming Your Wobble

Over the past decade we've helped all kinds of leaders—from seasoned politicians to Millennial entrepreneurs to the heads of multibillion-dollar companies—wrestle with trust issues. In doing so, we've learned a lot about strategies you can deploy to overcome your own trust

wobbles. Let’s explore what’s most effective for each of the drivers in our trust triangle.

Empathy

Most high-achieving leaders struggle with this one. Signaling a lack of empathy is a major barrier to empowerment leadership. If people think you care more about yourself than about others, they won’t trust you enough to lead them.

Empathy wobbles are common among people who are analytical and driven to learn. They often get impatient with those who aren’t similarly motivated or who take longer than they do to understand something. Additionally, the tools and experience of the modern workplace continually distract or prevent us from demonstrating empathy, by imposing 24-hour demands on our time and putting at our disposal all sorts of technologies that compete for our attention at any given moment. Our beeping and buzzing devices constantly assert our self-importance, sometimes smack in the middle of interactions with the very people we’re working to empower and lead.

We advise empathy wobblers to pay close attention to their behavior in group settings, particularly when other

people have the floor. Consider what often happens in a meeting: When it kicks off, most people feel very engaged. But as soon as empathy wobblers understand the concepts under discussion and have contributed their ideas, they lose interest. Their engagement plummets and remains low until the gathering (mercifully) comes to an end. Instead of paying attention, they often multitask, check their phones, engage in flamboyant displays of boredom—anything to make clear that this meeting is beneath them. Unfortunately, the cost of these indulgences is trust. If you signal that you matter more than everyone else, why should anyone trust the direction you're going in? What's in it for the rest of us to come along?

There's a basic solution to this problem. Instead of focusing on what you need in that meeting, work to ensure that everyone else gets what they need. Take radical responsibility for the others in the room. Share the burden of moving the dialogue forward, even if it's not your meeting. Search for the resonant examples that will bring the concepts to life, and don't disengage until everyone else in the room understands. This is almost impossible to do if texting or checking email is an option, so put away

your devices. Everyone knows you're not taking notes on their good ideas.

Indeed, the last thing we'll say on empathy is this: If you do nothing else to change your behavior, put away your phone more frequently. Put it truly away, out of sight and out of reach, not just flipped over for a few minutes at a time. You'll be amazed at the change in the quality of your interactions and your ability to build trust.

Logic

If people don't always have confidence in the rigor of your ideas, or if they don't have full faith in your ability to deliver on them, then logic is probably your wobble. If they don't trust your judgment, why would they want you at the wheel?

When logic is the problem, we advise going back to the data. Root the case you're making in sound evidence, speak about the things you know to be true beyond a reasonable doubt, and then—this is the hard part—stop there. One reason Larry Bird was such an extraordinary basketball player was that he only took shots he knew he could reliably make. That choice made him different from other great players who let ego and adrenaline cloud their shooting judgment. Bird studied and practiced so

relentlessly that by the time the ball left his hands in the heat of competition, he knew exactly where it was going. If logic is your wobble, take Bird's example and learn to "play within yourself."

Once you get comfortable with how that feels, start expanding what you know. Along the way, make an effort to learn from other people. Their insight is among your most valuable resources, but to access it, you must be willing to reveal that you don't have all the answers—something leaders often resist. Engaging people about their experience has the additional benefit of communicating who you are and what energizes you professionally—an authenticity boost.

For most logic wobblers, however, rigor isn't the issue. Much of the time, the problem is the perception of wobbly logic rather than the reality of it. Why does this happen? Because they're not communicating their ideas effectively.

There are generally two ways to communicate complex thoughts. The first takes your audience on a journey, with twists and turns and context and dramatic tension, until they eventually get to the payoff. Many of the world's best storytellers use this technique. You can visualize this approach by imagining an inverted triangle. The

journeying storyteller starts at the top, at the inverted base of the triangle, and traces an enchantingly meandering route down to its point.

If logic is your wobble, however, that's a risky path to take. With all that circuitous journeying, you're likely to lose your audience along the way rather than build trust in your judgment. Listeners may even abandon you at one of your narrative turns.

To avoid that, try flipping the imaginary triangle upright. Start with your main point, or headline, at the top of the triangle, and then work your way down, building a base of reinforcing evidence. This approach signals a clarity of vision and a full command of the facts. Everyone has a much better chance of following your logic. Even if you get interrupted along the way, you'll at least have had a chance to communicate your key idea.

Authenticity

If people feel they're not getting access to the "real" you—to a full and complete accounting of what you know, think, and feel—then you probably have an authenticity wobble.

A quick test: How different is your professional persona from the one that shows up around family and friends? If there's a sharp difference, what are you getting in return

for masking or minimizing certain parts of yourself? What's the payoff?

Being your "real self" sounds nice in theory, but there can be powerful reasons for holding back certain truths. The calculation can be highly practical at times, if wrenching—as in deciding to stay closeted in a workplace that's hostile to queer identities. There may also be times when expressing your authentic feelings may risk harmful consequences: Women, for example, are disproportionately penalized for displaying negative emotions in the workplace, and Black men are burdened by the false stereotype that they are predisposed to anger. We're not talking here about moments of prudent self-censorship, which sometimes can't be divorced from a larger context of bias or low psychological safety. Instead, we're talking about inauthenticity as a strategy, a way of navigating the workplace. If this is how you operate, you're dealing with an authenticity wobble.

In our experience, although withholding your true self may sometimes help you solve problems in the short term, it puts an artificial cap on trust and, by extension, on your ability to lead. When people sense that you're concealing the truth or being less than authentic, they're far less

willing to make themselves vulnerable to you in the ways that leadership demands.

We've observed the cost of inauthenticity up close in the performance of diverse teams. Diversity can be a tremendous asset in today's marketplace, and the companies that get it right often enjoy powerful competitive tailwinds. But this advantage isn't automatic. Simply populating your team with diverse perspectives and experiences doesn't always translate into better performance. In fact, the uncomfortable truth is that diverse teams can underperform homogenous teams if they're not managed actively for differences among members. That is due in part to a phenomenon called the common information effect, which works like this: As human beings, we tend to focus on the things we have in common with other people. We tend to seek out and affirm our shared knowledge, because it confirms our value and kinship with the group. Diverse teams, by definition, have less common information readily available to them to use in collective decision-making.

Consider two teams of three people, one in which the three members are different from one another, and the other in which they're similar. If both teams are managed

in exactly the same way—if they simply follow the same best practices in group facilitation, for example—the homogenous team is likely to perform better. No amount of feedback or number of trust falls can overcome the strength of the common information effect.

But the effect only holds if people wobble on authenticity. When they choose to bring their unique selves to the table—that is, the parts of themselves that are different from other people—they can create an unbeatable advantage by expanding the amount of information the team can access. The result is an inclusive team that's likely to outperform (by a long shot) both homogenous teams and diverse teams that aren't actively managed for inclusion. (See the exhibit "Trust, diversity, and team performance.")

Trust, diversity, and team performance

Diversity doesn't automatically confer advantages in decision-making. In fact, if diverse teams aren't managed actively for inclusion, they can underperform homogenous ones. That's because shared knowledge is key in decision-making, and diverse teams, by definition, start out with less of it. But if you create conditions of trust that allow diverse team members to bring their unique perspectives and experiences to the table, you can expand the amount of knowledge your team can access—and create an unbeatable advantage.

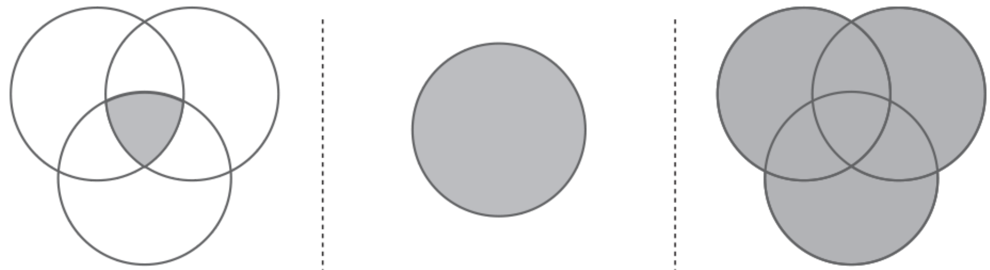

This expansion of knowledge and its obvious benefits rely on the courage of authenticity wobblers. We know how difficult sharing who we really are can be, and we also know that it's sometimes too much to ask. But if we regularly give in to the pressure to hold back our unique selves, then we suppress the most valuable parts of ourselves. Not only do we end up concealing the very thing the world needs most from us—our differences—but we also make it harder for people to trust us as leaders.

Here's the reason to care, even if you don't see yourself as different: All of us pay the price of inauthentic interactions, and all of us have a better chance of thriving in inclusive environments where authenticity can flourish. Gender bias, in other words, is not just a woman's problem. Systemic racism is not just an African American or Latinx problem. It's our shared moral and organizational

imperative to create workplaces where the burdens of being different are shouldered by all of us. After all, we will all benefit wildly from eliminating them.

One of the lessons we've learned in our work with organizations is that creating spaces where authenticity can thrive is not as hard as it may seem. It is an urgent, achievable goal that requires far less audacity than disrupting industries or growing complex organizations—things leaders do every day with deep conviction in the outcomes. If all of us take responsibility for creating companies where difference can thrive, and all of us take responsibility for showing up in them authentically, then our chances of achieving true inclusion—and building high levels of trust—start to look pretty good.

So pay less attention to what you think people want to hear and more attention to what you need to say to them. Reveal your full humanity to the world, regardless of what your critics say. And while you're at it, take exquisite care of people who are different from you, confident in the knowledge that their difference is the very thing that could unleash your potential and your organization's.

In Myself I Trust

We've argued that the foundation of empowerment leadership is getting other people to trust you. That's certainly true, but there's one last thing you need to know. The path to empowerment leadership doesn't begin when other people start to trust you. It begins when you start to trust yourself.

To be a truly empowering leader, you need to take stock of where you wobble not only in your relationships with others but also in your relationship with yourself. Are you being honest with yourself about your ambitions, or are you ignoring what really excites and inspires you? If you're hiding something from yourself, you've got an authenticity problem you need to address. Do you acknowledge your own needs and attend properly to them? If not, you've got to adopt a more empathetic posture toward yourself. Do you lack conviction in your own ideas and ability to perform? If so, you've got some logic issues to work out.

Doing this work is important as a leader, for an arguably obvious reason. If you don't trust yourself, why should anybody else trust you?

A Campaign to Rebuild Trust

Let's now return to Uber. When we began working with the company, it was certainly wobbling—so much so that we diagnosed it as "a hot mess."

What was going on?

Consider the basic trust-related facts. There's no question that Uber had empathy problems. The company's focus on growth at all costs meant that relationships with stakeholders, particularly drivers and employees, needed real attention. Riders also needed to be assured that their safety wouldn't come second to the company's financial performance. Additionally, despite its disruptive success, Uber hadn't answered questions about the long-term viability of its business model or about whether its managers had the skills to lead an organization of its expansive scale and scope. These were unaddressed logic problems. Finally, the company's war-room mentality was undermining its authenticity. In the "us versus them" culture at Uber, people were skeptical that they were getting the full story.

By the time Frances began working with Kalanick, he had already begun making changes to steady the company's trust wobbles. He had hired Eric Holder, for example, who had served as U.S. attorney general under President

Obama, to lead a rigorous internal investigation into harassment and discrimination—and when Holder made a sweeping set of recommendations, Kalanick took action to implement them. The company was also on the verge of rolling out new driver-tipping functionality, which would go on to generate $600 million in additional driver compensation in the first year of its launch. New safety features were in development, too, designed to give both drivers and riders additional tools to protect themselves.

Kalanick didn't get the chance to see most of these initiatives to completion, at least not from the CEO chair. In June 2017, he was forced out as CEO, although he retained his board seat and an equity stake in the company until December 2019, when he gave both up. He was ultimately replaced by Dara Khosrowshahi, the former Expedia CEO, who had a track record of effective leadership at the helm of young companies.

Frances soon began working with Khosrowshahi to continue the campaign to rebuild trust internally. Together they led an effort to rewrite the company's cultural values, one that invited input from all 15,000 employees on the principles that they wanted Uber to live by. The new motto they settled on was "We do the right thing. Period." Other

early trust wins for Khosrowshahi included strengthening relationships with regulators and executing a logic-driven focus on the services and markets that were most defensible.

Most of the work we did during this period was aimed at rebuilding trust at the employee level. Some things were easy to identify and fix, like ratcheting down the widespread, empathy-pulverizing practice of texting during meetings about the other people in the meeting, a tech-company norm that shocked us when we first experienced it. We introduced a new norm of turning off all personal technology and putting it away during meetings, which forced people to start making eye contact with their colleagues again.

Other challenges were harder to tackle, like the need to upskill thousands of managers. Our take was that Uber had underinvested in its people during its period of hypergrowth, leaving many managers unprepared for the increasing complexity of their jobs. We addressed this logic wobble with a massive infusion of executive education, using a virtual classroom to engage employees in live case discussions—our pedagogy of choice—whether they were in San Francisco, London, or Hyderabad. Although our

pilot program was voluntary and classes were sometimes scheduled at absurdly inconvenient times, 6,000 Uber employees based in more than 50 countries each participated in 24 hours of instruction over the course of 60 days. It was an extraordinary pace, scale, and absorption of management education.

The curriculum gave people tools and concepts to develop quickly as leaders while flipping a whole lot of upside-down communication triangles. Employees gained the skills not only to listen better but also to talk in ways that made it easier to collaborate across business units and geographies. Frances went out in the field, visiting key global offices in her first 30 days on the job, carving out protected spaces to listen to employees and communicate leadership's commitment to building a company worthy of its people. At a time when many employees were conflicted about their Uber affiliations, Frances made it a point to wear an Uber T-shirt every day until the entire company was proud to be on the payroll.

Within a year, Uber was less wobbly. There were still problems to be solved, but indicators such as employee sentiment, brand health, and driver compensation were all heading in the right direction, and the march toward an

IPO began in earnest. Good people were deciding to stay with the company, more good people were joining, and, in what had become our favorite indicator of progress, an increasing number of Uber T-shirts could now be spotted on city streets. It was all a testament to the talent, creativity, and commitment to learning at every level of the organization—and to the new foundation of trust that Kalanick and Khosrowshahi had been able to build.

Originally published in May–June 2020. Reprint R2003H

Cultural Innovation

by Douglas Holt

BUILDING THE NEXT billion-dollar innovation is an irresistible goal. To get a leg up, many companies now emulate the innovation model perfected in the tech sector. Procter & Gamble, for example, pursues what it calls *constructive disruption*. The company has designed its innovation process like a start-up's, with a venture lab that pulls in tech entrepreneurs and a lean probe-and-learn prototyping process.

That approach is not working. The reality is that in most consumer markets, innovation is a slow, incremental grind —extending master brands, adding a new bell or whistle, tweaking a formula. P&G's star innovations—such as a smart Pampers diaper that signals when a change is needed —aren't exactly threatening to become the next billion-dollar product.

And when companies do swing for the fences, they rarely achieve good results. Take Coca-Cola, which has long prioritized building a business in coffee. After years of research and testing, the company bet big on two innovations—Far Coast Coffee (a retail chain premised on sustainability) and Coca-Cola BlãK (Coke mixed with coffee). Both ideas failed badly, so the company eventually bought Costa Coffee, a British coffeehouse chain, at a steep price: $5 billion.

This problem is not an organizational one. Companies struggle because they put all their chips on one innovation paradigm—what I call *better mousetraps*. As Ralph Waldo Emerson noted long ago, "Build a better mousetrap, and the world will beat a path to your door." This is innovation as conceived by engineers and economists—a race to create the killer value proposition. It wins on functionality, convenience, reliability, price, or user experience. Better-mousetraps innovation is often the right bet if you're a tech company. Thousands of experts, seminars, and boot camps provide advice to help you on your way. But what about companies that operate in markets where new technology is less consequential or impossible to defend? For many of them, confronted with a pattern of poor return on

investment, chasing better mousetraps seems like an exhausting and expensive matter of running in place.

Fortunately, building better mousetraps is not the only way to innovate. In consumer markets, innovation often proceeds according to a logic I call *cultural innovation*. Think of Starbucks, Patagonia, Jack Daniel's, Ben & Jerry's, and Vitaminwater. Remember, innovation is in the eye of the beholder. When those brands broke through, consumers viewed them as major innovations, although a better-mousetraps perspective would reject that assessment. In each case people responded to the brand's ideology—a reimagining of the category that transformed the value proposition. Cultural innovations are embodied in distinctive products or services, to be sure, but also in founders' speeches, packaging, ingredients, retail design, media coverage, and even philanthropy.

The result? Those brands don't compete in the value-proposition race, trying to lead the category as it's currently defined; they play a different game. Better-mousetraps innovation is guided by quantitative ambitions: Outdo your competitors on existing notions of value. Cultural innovation operates according to

qualitative ambitions: Change the understanding of what is considered valuable.

I've spent the past 20 years researching and advising organizations on numerous cultural innovations. My work reveals the strategic principles that allow companies to pursue them—principles completely different from those used to build better mousetraps.

Ford Reinvents the Family Car

Buying a sport utility vehicle would have been an oddball idea for American middle-class families as late as 1989, but by 1995 the SUV was their unquestionable favorite, thanks largely to the Explorer—the pioneering vehicle that earned Ford roughly $30 billion in operating profit over its first decade. A spartan enclosed truck, the Explorer was yanked from its traditional role as functional transport on farms and ranches to become the aspirational choice of suburban families for commuting, delivering youngsters to school, and heading out to the mall. It succeeded wildly despite violating the rules of better mousetraps at every turn. It was a classic cultural innovation, targeting a fatal flaw in the family car culture of that era.

Idea in Brief

The Context

Most companies take a “better mousetraps” approach to innovation, improving a product’s functionality—with only average results.

A Different Approach

A few take a cultural innovation approach instead, first identifying a weakness in the existing category and then reinventing the category’s ideology and symbolism.

The Results

The Ford Explorer, for example, replaced the boring “mom mobile” minivan as America’s favorite family car with a promise of excitement, adventure, and glamour—even though the SUV wasn’t a technically superior vehicle.

The modern station wagon was a staple of the postwar nuclear-family ideal. All the major makes and models competed within this culture of suburban functionality. In the 1980s minivans rapidly replaced station wagons, winning on important benefits—plenty of seats, great storage, easy entrance and egress—that allowed families to haul kids and their friends around town and on summer trips.

The minivan's pragmatic design and ubiquity created a big symbolic problem. Vehicles are judged as much for the identity they project as for function: Status, sophistication, and masculinity all play a role in creating "premium" cars, which at the time were predominantly imports. Minivans came to represent the quotidian life of suburban parents, mocked as the centerpiece of a boring existence organized by "mom mobile" routines. Parents began to yearn for a car that would replace this stigma with an aspirational identity.

In the 1980s the Reagan-era revival of America's frontier ideology, which championed rugged individualists taming wild nature, inspired a critical mass of urban and suburban residents to reimagine the family car as a swashbuckling vehicle for off-road adventures. The offerings at the time were a poor fit for families: The Jeep Cherokee (XJ) and Chevy's massive Suburban were rough-driving trucks that lacked the amentities of passenger cars. The Ford Bronco and the Chevy S-10 Blazer offered only two doors. Nonetheless, many families were willing to forgo the minivan's creature comforts for the symbolic value that trucks bestowed. It took the incumbent automakers the better part of a decade to engage with this opportunity.

They were lucky to be in an industry with very high barriers to entry; otherwise they would no doubt have been beaten to market by a challenger brand.

Eventually the big three domestic truck players—Ford, General Motors, and Chrysler Jeep—raced to bring a comfortable, luxuriously equipped four-door SUV to market. The winner would be the brand that managed to seduce parents into thinking about family cars in a new way. Jeep had the initial advantage, given its potent off-road pedigree, and its new Grand Cherokee, launched soon after the Explorer, won many plaudits. However, Jeep's idea of a family SUV was a straight take on the frontier-adventure myth, showcasing performance on wilderness outings—a myth better aimed at young single men than at upscale families.

The Explorer was launched with advertisements that dramatized a new ideal of family life, rejecting the dull suburban minivan. Ford made two crucial changes to the frontier-adventure myth, both of which connected powerfully with parents. Instead of Jeep's macho excursions, the company offered a vision of families communing in the wilderness. Ads showed them whisking off to remote places in an Explorer to make memories while

gathering under the stars, kids happily trading in their tech for spiritual contentment. And parents who owned an Explorer got to have a life too. Ads showed them escaping on urban adventures—eating at boutique restaurants or attending the theater. They might live in the suburbs, but they could still enjoy a cosmopolitan life.

Families flocked to the Explorer. Sure, most of the time they were still hauling groceries and dropping kids off at soccer practice, just as they would have done with a minivan. But they were buying into a myth. Driving an Explorer allowed them to feel they'd finally escaped the world of mom mobiles for a more adventurous life.

In the postwar era, safety was a modest concern, despite Ralph Nader's best efforts. Even getting people to use seat belts was a challenge. By the early 1990s, though, car safety had captured the public's imagination owing to two big better-mousetraps innovations—airbags and antilock brakes—that were promoted heavily in auto advertising and the media.

Ford discovered early on that people believed that the huge size and weight of SUVs made them uniquely safe and that their off-road capabilities meant they were especially skid-resistant in bad weather. So the company crafted a

sales pitch to reinforce that perception. The car's elevated seats conferred a feeling of power and invincibility, particularly for women. When couples came to a dealership, the salesperson would ask the woman to test-drive the Explorer so that she could appreciate the feeling of safety from the high perch. Ford was able to persuade customers that they were buying the safest car on the road.

The Explorer was a great success, comparable to celebrated Silicon Valley innovations in terms of its market impact and profitability. Yet its breakthrough is incomprehensible when viewed through the lens of better mousetraps. The vehicle was not an engineering advance—quite the opposite. It relied on dated technology. Explorers accelerated lethargically. They were top-heavy and cornered poorly. They cost a lot and were far more expensive to maintain than minivans. And they were gas-guzzlers that generated enormous increases in CO_2. But families were willing to pay near-luxury prices because the SUV perfectly addressed the symbolic problem in the market's status quo.

The Cultural Innovation Model

Let’s look at a second case—Blue Buffalo dog food—to recognize the key steps in cultural innovation and to explain why incumbents often fail at it.

For decades Nestlé Purina, Mars, and Procter & Gamble dominated the profitable U.S. dog food category with powerful brands, distribution muscle, strong R&D, and big marketing budgets. Yet all three were beaten badly by Blue Buffalo, a tiny start-up, which was so successful that General Mills eventually bought it for $8 billion, while Procter & Gamble threw in the towel and sold its entire pet food division to Mars for less than $3 billion. Blue Buffalo bested the established brands by reinventing dog food culture. Here’s how.

Step 1. Deconstruct the category’s culture

Markets are belief systems embraced by those who participate in a category: companies, consumers, and the media. To understand your category’s culture, think like a sociologist. Step back and make the familiar strange. What are the category’s taken-for-granted organizing principles? What is the dominant ideology?

Before Purina launched the modern industrial dog food category, in the 1920s, most American families fed their dogs table scraps. Purina’s standardized extruded kibble

made inroads with consumers, and by the postwar era the company had adopted the mass-marketing techniques pioneered by food manufacturers such as Kraft and General Mills. Its ads featured heart-tugging images of cute dogs and their loving owners. The implicit message was "Purina is the biggest, best-known dog food company, so of course you can trust us to make food that will keep your dog healthy and energetic." Ingredients were rarely mentioned.

The category's first cultural innovation came in the 1970s, on the heels of media hype about scientific findings that certain vitamins and superfoods could keep people healthy. (Fiber and antioxidants were hot topics.) Cultural innovators, led by Hill's Science Diet and Iams, championed a new, scientific dog food ideology. The companies produced separate products for the various stages of a dog's life. Marketing featured veterinarians announcing cutting-edge formulas based on the best nutritional science. These products were sold in vets' offices—the ultimate sign of medical credibility. Purina launched a fast-follower grocery brand, Purina ONE, with ads featuring scientists in lab coats and packaging full of medical terminology.

These new brands taught owners to value dog food primarily for its nutritional benefits and offered them a scientific lexicon that "proved" quality nutrition. They encouraged owners to view the making of pet food as a complex scientific endeavor. The ingredients, however, remained hidden in small print.

Step 2. Identify the Achilles' heel

Categories' cultures eventually develop a fatal flaw, and cultural innovators pinpoint the emerging vulnerabilities. Throughout the early 2000s America's industrial-scientific food culture was subject to damning critiques in the media and by dozens of insurgent anti-industrial food movements. Dog owners began to feel similar concerns; they questioned whether those bags of kibble made by big companies were actually good for their pets. Then, in 2007, thousands of dogs and cats died after eating contaminated pet food. The media reported that one ingredient, wheat gluten contaminated with melamine, was bulk-sourced from China. Owners had had no idea that they were feeding their dogs wheat gluten or that it was imported from China. They began to take far more interest in the actual ingredients of dog food.

Step 3. Mine the cultural vanguard

Category transformations are usually prefigured by ideas and practices worked out at the margins. When cracks form in a category's culture, a *cultural vanguard* often appears before big companies show up. Innovators study the vanguard closely, and even participate in it, to find a strategic direction for their challenger ideology and the symbols required to bring it to life.

A small "natural" dog food subculture, separate from the national brands, had developed in prior decades. Alternative-health companies and their avid customers believed that healthful dog food should emulate what dogs ate before they became domesticated. The subculture's brands, which were sold in boutiques and natural-foods stores, were very expensive and marketed to niche customers. They made little effort to win converts from the big industrial-scientific brands.

A cultural innovation framework

Blue Buffalo upended industry giants like Purina and P&G by reconfiguring the category's ideology, using potent symbols. As a result, it transformed the value proposition for dog food.

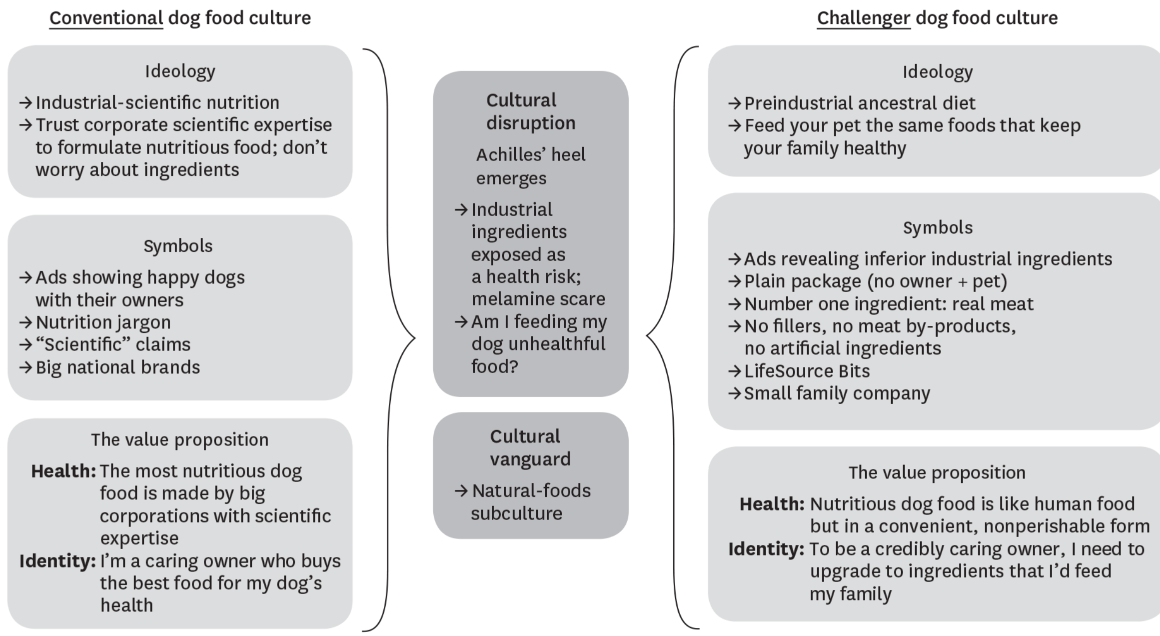

The brands lionized whole ingredients and transparent supply chains. They were all about real meat, poultry, and fish, along with whole-food carbohydrates (sweet potatoes, rice), and they fastidiously avoided anything artificial. The subculture encouraged customers to beware of "fillers" (processed starches such as corn, wheat, and soy) and meat by-products. Their packaging highlighted ingredients rather than happy dogs and loving owners.

Step 4. Create an ideology that challenges the Achilles' heel

Cultural innovators source materials from the vanguard to build a new brand concept. The natural-foods subculture's ideology was hidden: Alternative-health zealots talked to

one another and used rhetoric aimed at the already converted. Blue Buffalo, which was founded in 2002 by a Connecticut family that had become obsessed with the link between pet diet and health after their Airedale terrier (named Blue) died of cancer, acted as the subculture's proselytizer. The brand challenged the weak assumption that anchored the industrial-scientific ideology—that kibble was surely nutritious, even though owners had no idea what the compressed brown pellets were made from. In doing so, it created a litmus test for responsible dog ownership.

Blue Buffalo pushed owners to evaluate dog food as *food*. Those other kibble brands were full of industrial products that pet owners would never eat. People needed to take control and make sure their dog food contained healthful ingredients, no different from what they'd feed their families. Blue Buffalo's pet food was made with the same ingredients as a good human diet, so by switching brands, owners could ditch their newfound guilt and claim an enlightened identity—they really did feed their dogs nutritious food.

Step 5. Showcase symbols that dramatize the ideology

Cultural innovations are brought to life by a combination of symbols that dramatize them in the most compelling manner. They select symbols from the marketing mix that work together, attack the Achilles' heel, and draw a clear contrast with the category's dominant culture.

Blue Buffalo leveraged the leading symbols of the natural-foods subculture and created additional symbols to illustrate the notion that Blue Buffalo was, in effect, the same healthful food that owners themselves ate, converted into a compact, convenient, nonperishable form. The company repurposed the subculture's four foundational claims—real meat is the number one ingredient, no meat by-products, no fillers, nothing artificial—and used them in dozens of low-budget ads, produced to look like documentaries: Owners gathered in a living room, comparing notes on their preferred dog foods. Some were taken aback to read that their favorite brand contained "chicken by-product," while Blue Buffalo users proudly proclaimed that the first ingredient in theirs was deboned chicken. The company taught owners to read the label the next time they considered buying a bag of kibble.

And Blue Buffalo developed its own mini-kibble: LifeSource Bits—small, dark-purple (rather than brown)

balls made with superfoods such as blueberries, flaxseed, cranberries, and kelp. The company pushed owners to draw a connection between what their families ate to avoid chronic disease and what would give their dogs the same kind of protection.

As Blue Buffalo's challenge worked its magic, millions of owners decided to spend far more on dog food to avoid guilt. They bought into an entirely new value proposition: a new nutritional benefit (healthful dog food contains the same ingredients that healthful human food does) and a new identity benefit (switching to Blue Buffalo proved that they were truly caring owners).

Why Incumbent Counterattacks Failed

Despite the company's strategic brilliance, Blue Buffalo should never have been able to build a business that was worth $8 billion. The three incumbents completely dominated the market and should have prevailed over the upstart. All three invested heavily in new brands and line extensions, but they struck out because, working with a better-mousetraps mindset, they misunderstood the nature of Blue Buffalo's cultural innovation.

Iams: Cultural incoherence

P&G believed that Blue Buffalo was gaining ground by making a big deal of a simple "new and improved" ingredients claim. The company assumed that if it matched those ingredients with a line extension, owners would choose the trusted brand over Blue Buffalo. So P&G launched Iams Healthy Naturals, featuring two of Blue Buffalo's ingredients claims (no fillers, no artificial ingredients), with a big ad campaign and promotions. When that attempt failed, the company tried a more expensive iteration, Iams Naturals, which had meat as the number one ingredient. But to no avail.

What went wrong? Both products relied on brand names that tried to knit together the dominant industrial-scientific ideology (which Iams had championed for decades) with the natural dog food subculture—and the result was culturally incoherent. Iams came off as an impostor. It didn't help that the company's advertising campaigns used exactly the same trope (loving owner playing with energetic pet) that industrial-scientific brands had relied on for 40 years instead of showcasing ingredients, a key concern in the natural pet food subculture. P&G unwittingly sabotaged its rebuttal with its confused symbolism.

Purina: Purpose gone awry

Purina, too, launched a line extension—Purina ONE Beyond—to defend against Blue Buffalo. The effort led with not one but two industrial-scientific brand names (Purina and ONE), inadvertently signaling to consumers that this was not a credible natural dog food.

In addition, the company (which fancied purpose-driven branding at the time) decided to tie Beyond to a purpose. It knew from trends research that upscale owners favored green products, so it decided that Beyond would be the dog food that helped save the planet. An anthemic launch ad, depicting a glowing field, proclaimed, "We believe together we can make the world a better place one pet at a time." The problem was that environmental sustainability had nothing to do with Blue Buffalo's challenge, which centered on nutrition and health. Dog owners simply ignored Beyond.

Mars: A mismanaged acquisition

Incumbents' standard response when threatened by cultural innovation is to buy the threatening company or a close competitor. In 2007 Mars did just that by acquiring Nutro, a strong brand in the natural pet food subculture and a credible challenger to Blue Buffalo. That was a

promising move. To make it work, though, Mars would have had to shift Nutro marketing to attack industrial dog food, copying Blue Buffalo. It's unlikely that Mars ever considered that move, which would have meant attacking its biggest brand, Pedigree. Instead managers did just the opposite: They converted Nutro to a mass-marketing approach using ads little different from those of Iams.

P&G, Purina, and Mars never understood that they were fighting an existential battle to sustain their brands' authority as experts on healthful, nutritious dog food—not just racing to clean up their ingredients panels. As a result, Blue Buffalo convinced millions of dog owners that a product once viewed as a fussy extravagance was actually a necessity for people who truly loved their dogs.

Stuck in the Better-Mousetraps Mindset

Cultural innovation has often been an entrepreneur's gambit. Even when incumbents happen upon extraordinary cultural opportunities that should be easy to spot and straightforward to execute on, they fail time and again. If companies are to succeed at cultural innovation, they need to avoid three pitfalls.

Working eternally in the present

Even if they don’t think in such terms, companies are masters of their category’s existing culture. They have to be to excel at their current business. Their metrics and planning focus on it. As a result, managers come to perceive the category as an immutable reality, even though it’s actually built on a fragile consensus. If you’re trapped in the present tense, it is extremely difficult to examine the category from the outside and identify its emerging flaws. These ideological blinders explain why hundreds of highly trained professionals at the biggest pet food companies responded inadequately when Blue Buffalo attacked their billion-dollar businesses.

Being wedded to a product’s features

The better-mousetraps paradigm assumes that a product’s features are objective characteristics that consumers value. As a result, products are construed in building-block terms—as stacks of features that together create a value proposition. Innovation, then, requires improvements to particular features that consumers value. But features aren’t just building blocks— they can be malleable cultural symbols of an ideology. The incumbent dog food companies assumed that Blue Buffalo was simply offering

trendy new ingredients claims. But in fact those claims became “evidence” in Blue Buffalo’s whistleblower project, revealing that owners had been hoodwinked by the industrial-scientific brands.

Ignoring the value of identity

The better-mousetraps paradigm views innovations as great functional achievements, but that overlooks a critical component of many innovations: bolstering aspects of consumers’ identity. Ford, as we have seen, persuaded customers that they could trade in the dreary suburban minivan lifestyle for outdoor adventure and sophisticated city excursions. Blue Buffalo consumers traded up to garner status as enlightened dog owners.

In 1995 Clay Christensen introduced one of the most influential ideas in business: disruptive innovation. He famously asked why great companies fail when they’re doing everything right. Christensen’s answer: Incumbents focus on serving the most-demanding customers with the best products because margins are high. So entrants provide simple, cheap, “underperforming” solutions to low-end niches. Incumbents tend to ignore segments with

poor margins and “inferior” products until it’s too late. If one were to turn Christensen’s advice into a mantra, it might be “Think like a cheapskate.”

But that’s not the only innovator’s dilemma. Great companies are also disrupted by innovations that don’t involve new technologies; a cheap, low-performance product; or a price-sensitive target. Incumbents are so intent on winning the category as it’s currently defined that they fail to identify cracks in its foundation. Cultural innovators outmaneuver them because they look for opportunities to blow up the dominant ideology in favor of a new regime. So for incumbents to innovate, they’ll need to adopt a second mantra: “Think like a cultural entrepreneur.”

Originally published in September–October 2020. Reprint R2005J

The Rules of Co-opetition

by Adam Brandenburger and Barry Nalebuff

THE MOON LANDING just over 50 years ago is remembered as the culmination of a fierce competition between the United States and the USSR. But in fact, space exploration almost started with cooperation. President Kennedy proposed a joint mission to the moon when he met with Khrushchev in 1961 and again when he addressed the United Nations in 1963. It never came to pass, but in 1975 the Cold War rivals began working together on Apollo-Soyuz, and by 1998 the jointly managed International Space Station had ushered in an era of collaboration. Today a number of countries are trying to achieve a presence on the moon, and again there are calls for them to team up. Even the hypercompetitive Jeff Bezos and Elon Musk once met to discuss combining their Blue Origin and SpaceX ventures.

There is a name for the mix of competition and cooperation: *co-opetition.* In 1996, when we wrote a book about this phenomenon in business, instances of it were relatively rare. Now the practice is common in a wide range of industries, having been adopted by rivals such as Apple and Samsung, DHL and UPS, Ford and GM, and Google and Yahoo.

There are many reasons for competitors to cooperate. At the simplest level, it can be a way to save costs and avoid duplication of effort. If a project is too big or too risky for one company to manage, collaboration may be the only option. In other cases one party is better at doing A while the other is better at B, and they can trade skills. And even if one party is better at A and the other has no better B to offer, it may still make sense to share A at the right price.

Co-opetition raises strategic questions, however. How will the competitive dynamics in your industry change if you cooperate—or if you don't? Will you be able to safeguard your most valuable assets? Careful analysis is required. In this article we'll provide a practical framework for thinking through the decision to cooperate with rivals.

What Is Likely to Happen If You Don't Cooperate?

If a cooperative opportunity is on the table, start by imagining what each party will do if it's *not* taken. What alternative agreements might the other side make, and what alternatives might you pursue? If you don't agree to the deal, will someone else take your place in it? In particular, will the status quo still be an option?

Let's start with a simple example. Honest Tea (which one of us cofounded) was approached by Safeway supermarkets to make a private-label line of organic teas. The new line would undoubtedly eat into Honest Tea's existing Safeway sales. So even though the supermarket was offering a fair price, the deal would ultimately be unprofitable for Honest Tea.

However, if Honest Tea didn't cooperate, Safeway would surely find another supplier, such as rival tea maker Tazo. Honest figured that if it took the deal, it could design the new Safeway "O Organics" line to resemble the flavors and sweetness of Tazo's products and compete less against its own. If Honest had said no, Tazo would probably have said yes and targeted Honest's

flavors, leading to the worst possible outcome. So Honest agreed to the deal.

Yet the company turned down a similar request from Whole Foods because the grocery chain insisted that the private line include a clone of Moroccan Mint, Honest's best-selling tea at the time. Honest didn't want to compete so directly against itself and believed that its rivals would have trouble copying the tea—which indeed turned out to be true.

Idea in Brief

The Context

The idea that competitors should sometimes cooperate with one another has continued to gain traction since it was initially explored in the 1990s.

The Issue

Even so, executives who aren't comfortable with "co-opetition" bypass promising opportunities.

A Framework for Action

Start by analyzing what each party will do if it doesn't cooperate and how that decision will affect industry dynamics. Sometimes cooperation is a clear win. Even if it isn't, it may still be preferable to not cooperating. But it's

critical to try to figure out how to cooperate without losing your current advantages.

UPS had to think through a similar opportunity when DHL, which had acquired Airborne Express some years earlier and was suffering large losses, asked UPS to fly DHL's packages within the United States. UPS had the scale to make the service efficient (potentially saving DHL $1 billion a year) and was already providing a similar service to the U.S. Postal Service, so the opportunity appeared to be a profitable one that would allow UPS to rent out space on planes it was already flying.

That said, *not* cooperating might have been even more profitable in the long run. If DHL's continuing losses led to its exit, UPS stood to gain much of DHL's U.S. market share.

But if UPS turned the deal down, DHL might have offered it to FedEx. And if FedEx accepted it, DHL would still be in the market and UPS would have lost out on potential profits. So UPS agreed to DHL's proposal, announcing a deal in May 2008. (It turned out to be not enough to save DHL, which decided during the recession later that year to leave the market.)

In the tech industry, thinking through alternatives to a deal is complicated because companies have multiple relationships with one another. Samsung's decision about whether to sell Apple its new Super Retina edge-to-edge OLED screen for the iPhone X is a good example.

Samsung could have temporarily hurt Apple in the high-end smartphone market—where the Samsung Galaxy and iPhone compete—by not supplying its industry-leading screen. But Apple isn't the only rival Samsung has to worry about. In addition to being one of the world's largest phone manufacturers, Samsung is also one of the largest suppliers to phone manufacturers (including Apple, across several generations). If it hadn't provided its Super Retina display to Apple, Apple could have turned to LG (which supplies OLED screens for Google's Pixel 3 phones) or BOE (which supplies AMOLED screens for Huawei's Mate 20 Pro phones), strengthening one of Samsung's screen-technology competitors. Plus, Apple is well-known for helping its suppliers improve their quality. Cooperating with Apple meant that Samsung would get this benefit and that its screen-technology rivals would not. The fact that the deal would increase Samsung's scale and came with a big check

attached—an estimated $110 for each iPhone X sold—ultimately tilted the balance toward cooperating.

It takes two to cooperate. Now let's look at the deal from Apple's perspective. Would it make Samsung a more formidable rival? It probably would: In the year prior to the iPhone X launch, revenue from Apple accounted for almost 30% of the Samsung display business, a division that generated $5 billion in profits. (Apple was also buying DRAM and NAND flash memory chips, batteries, ceramics, and radio-frequency-printed circuit boards from Samsung.) But for Apple, getting the best screen was worth bankrolling an already well-resourced rival—at least for a while.

The underlying economic reason that working together was advantageous to *both* sides was that Samsung had the best screen and Apple had a loyal customer base. Without cooperating, neither company could get the extra value from putting the superior screen on the new iPhone.

Will Cooperation Give Away Your Competitive Advantage?

Suppose you've analyzed the alternatives to cooperation and tentatively decided to move ahead. Doing so may mean sharing your special sauce. Then it might not be so special, and that could be a real problem. To get a read on the potential risk, figure out which of these four categories the deal falls into:

Neither party has a special sauce at risk, but the parties' combined ingredients create value

In this scenario neither side is giving anything away. A recent example is Apple and Google's decision to cooperate in creating contact-tracing technology for Covid-19. By sharing user location data across platforms, the two companies enabled governments and others to create effective notification apps. The circumstances here are exceptional, but it's not unusual for rivals to team up to set standards and create interoperability protocols and thereby create a bigger pie they can later fight over.

Both parties have a special sauce, and sharing puts them both ahead of their common rivals

In 2013, Ford and GM agreed to share transmission technologies. This made sense because they had

complementary capabilities: Ford led in 10-speed transmissions, GM in nine-speed. The arrangement saved both money, had no significant strategic impact, and freed their engineers to work on next-generation electric vehicles, giving each company a leg up on other automakers.

There's a caveat here: Cooperation is more challenging if the playing field isn't level at the start. GM turned down an opportunity to collaborate with Ford on a next-generation diesel engine for super-duty pickup trucks. Though the potential cost savings were compelling, Ford already had a competitive advantage in the F-150's lightweight all-aluminum body, and GM feared that without differentiation between engines, Ford would have an unbeatable edge.

Sometimes, getting ahead of (or not falling behind) other rivals outweighs considerations of relative advantage. Autonomous driving technology, for instance, will be a key capability in the near future. Most automakers recognize that they won't be able to develop self-driving vehicles quickly or cost-effectively alone. That's why Ford invited Volkswagen to join its investment in Argo AI, an autonomous vehicle start-up.

VW's $2.6 billion investment (along with its $500 million purchase of Ford's shares of the start-up) greatly reduced the drain on Ford's resources.

The deal also plays to each party's respective strength in getting regulatory approvals—Ford is strong in the United States, VW in Europe—significantly increasing the chance that Argo AI will be one of the platforms that gets worldwide approval. Ford also believed that if it didn't work with VW, VW would find another partner, which would decrease the chance that Argo AI would become one of the approved standards.

Because Ford's market share is greater than VW's in the United States and VW is ahead of Ford in Europe, it was a good bet that this partnership wouldn't change the balance of power between them. The focus was on elevating the pair relative to their many rivals.

One party has a strong competitive advantage, and sharing only heightens it; even so, less-powerful parties are willing to cooperate

Amazon gives rival sellers on Amazon Marketplace access to its customers and warehouses. Why? For starters, while it loses some direct business and the associated markup,

it makes a commission on Marketplace sales. The net effect on profit depends on how the commission compares with the markup, and whether Amazon Marketplace (which accounts for $50 billion of the company's revenue) leads to an increase in the company's total volume.

Even if the net effect were negative, blocking rival sellers from its platform would push them to other sites that could compete with Amazon. More important, though, when Amazon shares its platform, it becomes a hub—the starting place for any search. It makes money when a person looking for a book or a computer cable comes to its site and purchases additional, higher-margin products like electronics or clothing. Amazon also learns about the customer's preferences and can use this data to offer better recommendations and more accurately identify which Amazon-branded products to offer. And finally, opening up Amazon Marketplace allows Amazon to operate more warehouses and increase shipping volume, thereby reducing shipping times and lowering overall costs.

But why do other merchants cooperate with Amazon? Each partner, acting individually, finds it more profitable,

even necessary, to be part of the Amazon ecosystem. But it's a collective action problem: When the merchants all join its platform, they make Amazon a more formidable rival. Indeed, both the European Commission and the U.S. House Subcommittee on Antitrust, Commercial, and Administrative Law are investigating whether Amazon Marketplace is using its dominant position to undermine and compete unfairly with its merchant "partners."

One party shares its secret sauce to reach another's customer base, even though doing so carries risks for both parties

We saw this dynamic when Samsung shared its high-end screen with Apple. Google and Yahoo provide another example.

Google is better than any of its rivals at turning ads that appear alongside searches into clicks—that's its secret sauce. In 2008 it agreed to do ad placement for Yahoo. Google's technology would generate substantially more revenue per search for Yahoo, and sharing it was the quickest, surest way to extend its value to the market Google didn't already have. (In the short run, Google was unlikely to capture all of Yahoo's customers. By 2020,

Yahoo's share of search was down to 1.6%, but that decline took a dozen years.)

The potential gains were enormous. Given Yahoo's then 17% share of the $9 billion market, a 50% to 60% revenue increase would create almost $1 billion in annual profits to be split between the two companies.

The deal did carry some risks for Google. It might have made Yahoo into a stronger competitor, but that possibility was less worrisome because Yahoo was already cash-rich owing to its stake in Alibaba. (More cash probably wasn't material to its competitive position.) Improved ad technology on Yahoo might have led some Google users to switch, but it seemed unlikely that better ads would cause a large number to do so. Perhaps the greatest risk was that Yahoo would learn the recipe for Google's special sauce—but Google never planned to hand over its algorithms.

The risks for Yahoo were bigger. Its capabilities might wither if it became dependent on Google's black box. Were the partnership to end, Yahoo would be further behind, perhaps dangerously so. Those risks were mitigated by Yahoo's plan to continue doing ad

placement for its sites in Europe and thus maintain its own capabilities.

In the end the deal didn’t materialize; the U.S. Department of Justice ruled against it on the grounds that it might leave Yahoo a weaker competitor in the future. (One of us helped defend the agreement.) But the economics were compelling. One year later, Yahoo made a deal with Microsoft to have Bing provide its search ads.

It isn’t always possible to rent the sauce without giving away the recipe, however. Could the United States and China, for instance, cooperate on a mission to Mars? A seemingly insurmountable challenge is that it would involve sharing intellectual property that can’t be recaptured. This is a particularly sensitive issue since space technology spills over to military applications.

How to Structure an Agreement

The parties have almost gotten to yes. They’ve identified a desirable opportunity and found a way to share their special sauce without giving away the recipe. The remaining task is to craft the agreement. Two issues are particularly challenging when a prospective partner is

also a competitor: the scope of the deal and how the costs and benefits will be divided. (There may also be antitrust concerns; for more on those, see the sidebar “What About Antitrust Issues?”)

What About Antitrust Issues?

REGULATORS ARE NATURALLY SUSPICIOUS when rivals get together. Executives need to know which types of cooperation are permissible and which are not. Some antitrust violations are black-and-white: Businesses that coordinate to raise prices or divide up the market are engaged in collusion, pure and simple.

Regulators tend to take a more favorable view when businesses work together to reduce costs or expand demand. One good litmus test is to ask if customers will be better off as a result of the cooperation. For example, customers benefit if rivals partner to provide charging stations for electric cars. Similarly, supplying a rival tends to pass muster when it improves quality (as is the case when Samsung sells its Super Retina screens to Apple) and doesn’t foreclose market entry to other players.

There is always the possibility that regulators will step in to nix a deal, as they did with Yahoo’s 2008 agreement to have Google provide it with search ads. This is one of the challenges of co-opetition.

Establishing scope and control

First the parties have to figure out how far to extend their cooperation, who is in charge, and how they might unwind their arrangement should it no longer make sense.

The simplest types of cooperation are limited and don't raise control issues. In some cases one party becomes a nonessential supplier to the other—as Honest Tea did with Safeway or as CBS did when it supplied the show *Dead to Me* to Netflix. In other cases the parties share costs but not proprietary knowledge. Rival television stations sometimes share camera crews, for instance, and rival breweries coordinate on recycling. Several museums in a city may run an ad campaign or develop an all-access museum pass together. Generally these arrangements are easy to negotiate and can be unwound easily.

Agreements become challenging when one party has to cede control, however. Ford and GM's plan to share transmission technologies worked well at the R&D stage, but neither company was willing to give control of manufacturing to the other or even to a joint entity. Ford and GM could have written a contingent contract about who got what transmission production capacity when,

but this would have been tricky since demand is variable and transmissions are mission-critical. Fortunately, the majority of the cost savings came from using common designs and common parts, so Ford and GM limited the agreement to those areas.

In other circumstances one party is in charge and the other party is protected by a contingent contract with performance guarantees and penalties for not hitting specific targets. This works well in situations where there are established performance benchmarks. The party in charge, the one providing the guarantees, doesn't have to be told what to prioritize; instead the right-sized penalties allow it to internalize decisions and make calls that optimize the combined outcome.

It's important to structure any agreement in such a way that one side doesn't become dependent on the other. Otherwise, the dependent party may be backed into a corner when it comes time to renegotiate the deal—or distressed when the deal ends. As noted earlier, this was one of the Justice Department's issues with the 2008 Google-Yahoo deal.

Dividing the pie

Cooperation is an overall win-win, but splitting the gains is a zero-sum game. The solution is relatively straightforward when there's an even trade, as when Ford and GM shared transmissions. It's harder if cooperation involves an uneven trade and payments are required.

Consider interairline agreements to help stranded passengers. For a long time it was customary for airlines to take care of one another's passengers in the event of a flight cancellation, or what the industry calls an IROP (irregular operation). Airlines paid a low IROP rate to secure a seat on another carrier.

Cooperation broke down in 2015 when Delta thought other airlines were getting the better end of the deal and proposed a steep increase in the IROP rate. Delta was taking five American Airlines passengers for each Delta passenger that American took. American declined to pay more, and the agreement ended.

The underlying problem was an uneven trade. With an even balance of trade, the IROP fare doesn't matter. When the trade is out of balance, the right price is what ensures a fair deal. An IROP fare that was Delta's cost of a seat (including forgone sales to displaced customers) plus half the value of American's gains (the savings on a hotel

and meals and avoidance of the customer's ire) should have done the trick.

There might have been a way to save at least part of the deal without agreeing on price. Delta and American could have set up an agreement that guaranteed parity, trading seats on a one-for-one basis. If one airline had more cancellations and took more seats, the number of seats it got could be rationed going forward until things evened out.

The problem was ultimately resolved when the balance of trade was restored. After a series of computer outages and systemwide shutdowns, Delta found that it, too, needed some help. It renewed an agreement with American in 2018.

The challenges are greater when there are three or more parties to the deal and offsetting trades aren't possible. Take Ionity, a joint venture involving BMW, Daimler, Ford, Hyundai, Kia, and VW, which is building ultrafast electric-charging stations across Europe. The speed and cost-savings advantages from teaming up are enormous. Still, each partner has different geographical priorities, creating tensions over where to place the stations.

Splitting the massive price tag is even harder. It wouldn’t work to divide the costs equally; the partners have significantly different shares of the market, and Kia, with its much smaller slice, would walk away. Costs could be split according to market share—but should market share be based on unit sales, dollar sales, profits, or even miles driven? Each party had its favorite answer.

In the end the six companies agreed that costs would be divided in proportion to current unit sales. A simple, albeit somewhat arbitrary, heuristic like that may be a practical way to get a cooperative venture off the ground.

Changing Minds

Cooperation with rivals also has an important emotional aspect. Some people are comfortable with the idea that there can be multiple winners, and some are not. As a result, co-opetition may end up being a strategy of last resort even in cases where it should be a first resort.

Apple was on the verge of failure in August 1997 when Steve Jobs was finally forced to confront the fact that Microsoft was not the enemy. Jobs later admitted that “if the game was a zero-sum game where for Apple to win,

Microsoft had to lose, then Apple was going to lose." That change in perspective was hard for Apple loyalists to accept. When Jobs announced at the Macworld conference that Microsoft had invested $150 million in Apple, Bill Gates was booed.

Obvious opportunities for cooperation fall by the wayside when businesspeople don't focus on ensuring that all parties come out ahead. The world of check payments illustrates the problem.

Ever since printed checks were invented, more than 300 years ago, banks have needed a way to exchange those deposited by their account holders but written on other banks' accounts. The obvious solution was to establish a central clearinghouse. When the London banks failed to do this, the bank runners did it themselves. Instead of crisscrossing the city to exchange checks, they did an end run and all met at the Five Bells tavern. Some 50 years later the banks established the Bankers' Clearing House to do the same job.

In the modern era the U.S. Federal Reserve operated a system in which each bank would forward the paper checks it received to the Fed, which would then distribute them to the banks on which they were written. In 2001

some 40 billion checks were being flown around the country.

A logical alternative was to scan the checks and send digital images, thereby saving time and money. The challenge was that some of the small banks weren't set up to process digital images. Thus cooperation would further tilt the playing field. When the large banks didn't ensure that the small banks would also come out ahead, the small banks used their political power to block digital check clearing.

Then 9/11 forced the issue. With all planes grounded for over a week, checks were stranded and could not be cleared. At that point, the large banks finally agreed to ease the transition for small banks by having the Fed print the digital images and send the substitute checks to the small banks. In 2003 digital check clearing became established in law when Congress enacted the Check Clearing for the 21st Century Act.

It's also possible to work *around* mindsets. One solution is compartmentalization—both mental and actual. The Apple-Samsung deal, which happened during a billion-dollar legal battle between the two tech giants over patent infringements, was doubtless easier to arrange

given that Samsung operates as three separate companies with three separate CEOs. Apple could cooperate with one autonomous part of Samsung while competing with and suing another.

For a similar reason, we think it was wise for Ford to keep Argo AI, the autonomous vehicle start-up, a separate company. It was psychologically and contractually easier to get VW to invest in an entity that was outside Ford. The external structure helps ensure that the two will be equals and also makes it easier to bring in future partners.

Ultimately, getting the right mindset requires choosing the right people. The executives we interviewed emphasized the need to staff the cooperating teams with people who are open to the dual mindset of co-opetition.

That isn’t always easy, because people tend to think in either/or terms, as in either compete or cooperate, rather than compete *and* cooperate. Doing both at once requires mental flexibility; it doesn’t come naturally. But if you develop that flexibility and give the risks and rewards

careful consideration, you may well gain an edge over those stuck thinking only about competition.

We began this article with the missed opportunity for cooperation between the United States and the Soviet Union on a mission to the moon. Today the opportunities for countries to cooperate are even larger—from tackling Covid-19 and climate change to resolving trade wars. We hope that a better understanding of co-opetition will help businesses, managers, and countries find a better way to work and succeed together.

Originally published in January–February 2021. Reprint R2101C

Negotiating Your Next Job

by Hannah Riley Bowles and Bobbi Thomason

WHEN WE ASK PROFESSIONALS to describe a career negotiation, the first thing many people think of is bargaining with a hiring manager over an offer package. That scenario may spring to mind because compensation negotiations can be especially stressful and awkward and therefore become seared into our memories.

Although reaching agreement on pay and benefits is important, failure to think more broadly about your career could mean losing valuable opportunities for advancement. For instance, women are increasingly urged to negotiate for higher pay as a way to close the gender wage gap. However, studies have shown that women's "80 cents on the dollar" is explained more by

differences in men's and women's career trajectories than by differential pay for doing the exact same job. Our research and our work coaching executives suggest that negotiating your role (the scope of your authority and your developmental opportunities) is likely to benefit your career more than negotiating your pay and benefits does. And at times of work-life conflict, negotiating your workload and the conditions that affect it (including your responsibilities, your location, and travel requirements) may be critical to remaining gainfully employed and moving forward professionally.

As with other dealmaking, career negotiations should not be solely about getting as much as you can. The best negotiators generate mutually beneficial solutions through joint problem-solving and creative trade-offs, along with compromise. Furthermore, negotiating the direction of your career typically involves multiple stakeholders—including those in your personal life as well as those at work.

We advise professionals to think strategically about not just *what* they might negotiate but *how*. That means going beyond planning what to say at the bargaining table; it requires keeping your eye on larger objectives, ensuring

that you are negotiating with the right parties over the right issues, and preventing misunderstandings from derailing your requests or proposals because they are unconventional or potentially pathbreaking.

In the age of Covid-19, the time is ripe to improve your career negotiation skills. Many people are changing how they work (shifting to remote or flexible arrangements, for example), what they are working on (being redeployed or responding to new priorities), and with whom they're working (collaborating in new ways across functions and geographies). And transformations in our work lives are increasingly interlinked with transformations in our personal lives—whether that involves relocation decisions, periods of intense dedication to our jobs, or adapting to spikes in caregiving demands.

Drawing from a research project in which we collected thousands of stories from recent professional-school graduates, midlevel managers, and senior executives from seven global regions about how they advanced at pivotal points in their careers, we propose four steps for preparing for your career negotiations. They progress in a logical order, but you are likely to return to earlier steps

as your analysis proceeds. For instance, you might start out intending to negotiate for one type of opportunity but discover that you're better off negotiating for a different type. Or you might initially conceive of a proposal to present to your boss but then come to understand that your boss is actually not the key stakeholder who needs to be persuaded. Particularly for a complex and protracted negotiation, you should continually refine your analysis as you gain information.

1. Start with Your Career Goals

In our experience, negotiators too often start their preparation focused on the opportunity right in front of them, such as a job offer, rather than on their ultimate work and life aspirations. As you enter a period of change in your career, you should think about your short-and long-term aims and then map backward from those objectives to define the next steps you want to take. Don't forget to include quality-of-life considerations as well as professional ones. And be prepared to defer gratification if that's the right thing to do for the endgame.

Idea in Brief

The Problem

In job negotiations, professionals too often focus on pay and benefits and fail to think more broadly about how the opportunity will fit into their long-term career goals.

The Way to Start

Work backward from your career objectives to define the next steps you want to take. Be mindful of whether you're proposing something standard, an unusual arrangement for yourself, or an idea that will change the organization.

The Negotiation

Make sure you have all the necessary information and aren't operating under false assumptions. Explain why your request is also in the other party's interest. Seek input and feedback from people who could be helpful, and enlist allies to support your proposal.

Anya's story offers a cautionary tale. ("Anya" and all other individuals discussed in this piece are composites of case examples we studied.) When finishing her MBA program, she was evaluating two offers: one in consulting—the field she had previously worked in for several years—and one that would launch a new career in tech, which was what her heart truly desired. (Feeling torn between

two industries is common in job searches.) The consulting firm was offering her more money and status than the tech company was—unsurprising, given her track record in consulting and her limited experience in tech (one summer internship). Focused on the terms of the offers, Anya started her negotiation preparation wondering if she should walk away from the tech company unless it matched the salary offered by the consulting firm.

Making compensation the deciding factor can be a mistake. If we'd been coaching Anya, we would have encouraged her to start with her career goal: transitioning out of consulting into tech. We would have encouraged her to compare the competing offers not only with each other but also against her vision of what she wanted to achieve in her first five years out of graduate school. Next we might have asked, "To improve the tech offer, what might you negotiate to fulfill your dream of a career in tech?" After all, her lifetime earning potential could be higher in that booming sector than in consulting. Perhaps she could accept the lower compensation but negotiate for an accelerated promotion track—a solution that might appeal to the tech company because it would not need to

deviate from its compensation standards for MBA recruits.

Such longer-term thinking often pays off. In *Chasing Stars: The Myth of Talent and the Portability of Performance,* Boris Groysberg reports that the financial analysts who were most likely to retain their star status after moving to a new firm were those who had looked beyond pay and carefully researched whether the new firm would provide them with the organizational resources to excel. They understood that being successful in one setting doesn't guarantee success in another, so the compensation package was just one aspect of the job offer to consider. Our advice is to define from the start what you most want to achieve—whether that's being a top professional, making money, or living up to some other ideal—and then keep that goal in mind as your negotiation progresses.

2. Understand What You're Negotiating For

Career negotiations fall into three buckets. In *asking* negotiations, you propose something that's standard for someone in your role or at your level. In *bending*

negotiations, you request a personal exception or an unusual arrangement that runs counter to typical organizational practice or norms (for example, a remote work setup or a promotion to a position for which you lack the conventional qualifications). And in *shaping* negotiations, you propose ways to play a role in changing your organizational environment or creating a new initiative (such as revamping the way a project is run or launching a new business unit). Depending on whether you are in an asking, a bending, or a shaping negotiation, you will need to vary your arguments to win your counterparts' support.

In asking situations, you must demonstrate that your request or proposal is reasonable because it fits with existing practices or norms—for example, a pay raise is warranted in light of an outside offer, or you deserve a promotion or a developmental opportunity because other employees with your track record or experience have received such rewards. Asking negotiations often arise in the context of routine conversations about role assignments. If you are asked to do work that would move you away from your career goals, see if there is room for negotiation. For instance, you might be able to

explain why the proposed change in your role is not in the employer's interest: Perhaps it would hurt the performance of your team or damage the relationship with a high-value client. Another option is to agree to do the job for the sake of organizational needs in exchange for some other career-advancing opportunity. For example, you might say, "I will take on this role to help us out of the current crisis, but I would like to rotate into a job with more P&L responsibility after two years."

If you are in a bending negotiation—seeking some special exception or privilege—you need to keep your counterparts from doing what's easiest and simply saying, "No, that's not the way things are done around here." Justifying your request is particularly important if you are asking people to take a chance on you, such as putting you in a position for which you are not traditionally qualified.

Consider the case of Bela, who wanted to move from finance into a leadership role in IT as her company launched a digital transformation. The CIO considered her unqualified and seemed likely to dissuade the CEO from giving her the job. Bela came to realize that the CIO wanted someone more experienced to oversee the IT

transition, in part because failure would reflect poorly on the CIO's own leadership. So she asked for a six-month trial while the CIO searched for a potential replacement. Bela explained why her deep knowledge of the company's financial systems and her track record managing cross-functional teams prepared her to succeed in this IT role or, at a minimum, keep the company on solid footing until she was replaced by a new hire.

Although any negotiation can backfire, bending negotiations are particularly risky because they may give the impression that you're a prima donna seeking special treatment or unwilling to pay your dues. Deborah Kolb, an expert in career negotiations, suggests a role-play exercise to mitigate this risk: List the reasons why your counterparts would support your proposal; then come up with a list of reasons why they might say no anyway—and your possible responses. Beyond strategizing to get past "no," we advise weighing the downstream career risks and benefits of entering into an exceptional or unconventional work arrangement.

Whereas asking and bending negotiations are focused solely on your personal career path, shaping negotiations center on proposals to change the path of your

organization or working group. Because that commonly means seizing leadership opportunities, shaping negotiations typically involve more parties and the backing of allies.

Consider Samir's desire to lead a restructuring of his firm, which was run by an elite old guard that he saw as out of step with globally competitive business practices. Samir recognized that he needed to build a cross-generational coalition to support this change. As he made his case to key colleagues, he found allies among the veteran leaders who recognized that the firm's legacy would depend on retaining bold thinkers like him. He also found peers who appreciated his vision for growth. Finally, with his spouse's support, Samir worked out a plan to relocate internationally for another position if the firm rejected his proposal. He then began the negotiation process with confidence that he had enough buy-in within the firm to lead a transformational change, but also a satisfying alternative for himself and his family if that was not possible.

Organizations may be especially receptive to bending and shaping negotiations during challenging or fast-changing times, when people are looking for ways to

adapt and innovate. For instance, in light of the Covid-19 pandemic and social distancing restrictions, many employees need to change the way they work. Their collective bending negotiations are a useful source of information and experimentation for organizations and individuals trying to figure out how to maintain high morale and productivity during the crisis. Organizations are also welcoming shaping proposals from employees who have ideas about how to redeploy resources and open new markets in response to economic disruptions at home and abroad.

3. Reduce Ambiguity About What, How, and with Whom to Negotiate

No one would ever advise walking blind into a potential negotiation, but people do it all the time. One risk is that you'll "get Wahlberged," as the journalist Kate MacArthur put it, writing about how Mark Wahlberg negotiated a payment of $1.5 million to reshoot some scenes in a Hollywood film while his costar Michelle Williams accepted less than $1,000 for the same work. That case has been highlighted as an example of women's failure to

negotiate, but the underlying problem was a lack of information on what was negotiable. Williams had been led to believe that all the actors on the reshoot were effectively donating their time to save the film after another costar was pulled from the cast.

Reducing ambiguity is particularly important for ensuring that people from underrepresented groups—oftentimes women and people of color—get a fair shake. Many organizations are moving to make their recruitment and promotion practices more transparent so that all candidates have access to the same information and opportunities. Increasing transparency is obviously the responsibility of organizations, but individuals can take action too.

As you prepare to negotiate, write down all the questions you have: *What is potentially negotiable? How should I negotiate? Who will be my counterparts, and what do they care about?* There are many sources for this type of information. Talent professionals, for example, will explain in general terms what is typically negotiable and how (although they usually won't reveal the specifics of any individual case). Some information is available online. A media or YouTube search can give you

perspective on counterparts' points of view on strategic issues. A LinkedIn search can help you find professional contacts who may tell you more about a hiring manager or a department.

Although your personal and professional networks can be a valuable source of information, you should not rely on them alone to get an unbiased understanding of the situation. Think of a field in which men tend to be better paid than women. If women confer only with other women about customary salaries, and if men confer only with other men, women are likely to enter pay negotiations with lower expectations than men have—and to exit with worse outcomes.

Stretch your inquiry beyond your closest networks to ensure that you have the broadest information possible. Recently many people have been learning from how organizations in other industries or geographies are responding to the challenges presented by the Covid-19 pandemic. Better information helps generate innovative solutions; it can also help you make a persuasive case for managing your career the way you want to during these turbulent times.

4. Enhance Your Negotiations Through Relationships—and Vice Versa

As you aim to reduce ambiguity, you will undoubtedly think of people you might go to for information or advice. You might also think of others who could provide social support—those who would encourage and stand by you and give honest feedback if you are off track. Don't forget to identify potential advocates for your proposal. Who might be willing to speak up in favor of it? Who are your allies? Connecting with people who can be helpful is what we mean by enhancing your negotiations through relationships.

Consider the example of Brandon, an engineer who landed a job as a private equity associate after finishing business school. Lacking finance experience, he had been advised that his prospects of making partner were dim if he did not make a distinctive contribution. Brandon hoped to do that by arguing for the creation of a small fund to invest in marketable robotics projects—an underdeveloped growth area for the firm. Before negotiating to spearhead this initiative, he sought advice from his former robotics professor, who could spot

weaknesses in his proposal and help him fix them. He also found a partner at the firm who agreed to let Brandon shadow him on tech company boards.

To build a coalition of support for what you hope to do, you might start off by trying something akin to the shuttle diplomacy used by negotiators of international affairs: Make the rounds of key stakeholders, talking with them individually to solicit their feedback and input. Shuttling is more time-consuming than calling a summit of all interested parties (a meeting to pitch your proposal). But it enables you to privately explore people's interests and concerns and to incorporate their ideas into your game plan. It also helps you predict how people will respond when it comes time for you to present a formal proposal.

If you're concerned that shuttling around might make you appear conniving or manipulative, then be transparent about it. Explain that you're seeking input on an idea you have, and meet early with people who might block your proposal if they felt you weren't consulting them. To broaden buy-in, you might also enlist others to help you get feedback, keeping in mind Harry Truman's

words: "It is amazing what you can accomplish if you do not care who gets the credit."

Many of the negotiation cases we studied were rife with tales of conflict and resistance, but you needn't settle for compromises that leave both sides unhappy. The give-and-take that occurs when you're seeking a mutually beneficial deal can open your eyes to other perspectives, help you better understand your colleagues, and find ways of working together to create lasting solutions. In other words, career negotiations can enhance your working relationships—and we encourage you to strive for that outcome.

To generate goodwill and motivate agreement, we recommend that you explain to counterparts both why it is legitimate for you to be negotiating and how your proposal takes their interests into account. That's not always easy. For instance, we met one female executive who found out for the second time that a male subordinate was being paid more than she was. She probably wanted to say many things to senior leaders at her firm, but she chose the approach she knew would be most persuasive: "I know you are going to want to fix

this, because it is inconsistent with company practices and values."

Or take the example of Sandra, who ran the U.S. division of a major business unit and wanted to globalize it. To achieve her aim, she had to make a strategic case for why globalization was in the company's best interest and why she was the right person to lead the initiative. Addressing the hopes and concerns of managers both at headquarters and in the non-U.S. business units required numerous rounds of conversation in which she seeded and got feedback on her ideas. Sandra told us: "Over time, the logic [for globalization and my leadership] became compelling."

The four steps outlined above take time to implement—and there will be false starts and reversals. Most of the career negotiations recounted to us by senior executives, managers, and other professionals lasted weeks or months. They started with preliminary conversations that gradually evolved, particularly as new information or the entrance of new players influenced the way various

parties perceived their interests and the alternatives to agreement.

To maximize your odds of success, set targets for yourself that are specific and realistic—and that help hold you accountable to follow through with your plan amid pressing distractions and demands. Too often, negotiations fizzle or never get off the ground because larger goals become buried by everyday work.

One senior executive we interviewed told us, “You have a book to write of your life. Don’t let anyone else write your chapters.” We second that, but we also urge you to remember that great careers are not authored alone. Your narrative will be cowritten with work and life partners, and negotiation is at the heart of finding mutually gratifying ways for that story to unfold.

Originally published in January–February 2021. Reprint R2101E

Leading Through Anxiety

by Morra Aarons-Mele

THE CEO of a startup is sitting in the office space she recently leased for her fast-growing company. It's rush hour, but the streets outside are quiet, and so are the 600 empty cubicles outside her office door. Just yesterday her leadership team made the tough but crucial decision to send everyone home to work for the foreseeable future. In 30 minutes she needs to lead a videoconference to reassure her employees. But she's despondent, anxious, and just plain scared.

Versions of this scene have been playing out across the world over the past few months as Covid-19 cases rise and economies shut down. Founders, executives, managers, and employees have seen how fragile everything they've built has become—almost overnight. One evening back in March, my husband said to me, "I'm so scared, but I can't

let all the people who depend on me see that." He had been on hours of Zoom calls, trying to convince his staff and colleagues that they would get through the crisis. He was supposed to be the face of calm, but he was terrified.

How can you lead with authority and strength when you feel anxious? How can you inspire and motivate others when your mind and heart are racing? And if you hide the fear in an attempt to be leaderlike, where does it go?

Anxiety, of course, has a purpose. It protects us from harm. Psychologist Rollo May first wrote in 1977: "We are no longer prey to tigers and mastodons but to damage to our self-esteem, ostracism by our group, or the threat of losing out in the competitive struggle. The form of anxiety has changed, but the experience remains relatively the same." In other words, even though humans today aren't chased by predators, we are chased by uncertainty about the health of our loved ones, whether we'll have a job next week or next year, whether our company will go bankrupt—worries that provoke the same neurological and physical responses.

The Big Idea: Managing in an Anxious World

"LEADING THROUGH ANXIETY" is the lead article of **HBR's The Big Idea: Managing in an Anxious World**. You can read the rest of the series at hbr.org/anxiety:

- "When Anxiety Becomes Unbearable," by Gretchen Gavett
- "5 Ways Leaders Accidentally Stress Out Their Employees," by Tomas Chamorro-Premuzic
- "How History's Great Leaders Managed Anxiety," by Alison Beard
- "Anxiety as a Public Health Issue," by Sandro Galea
- "How We Experience Anxiety Today," by Kelsey Gripenstraw

According to the Anxiety and Depression Association of America, "Stress is a response to a threat in a situation. Anxiety is a reaction to the stress." Anxiety is fear of what might happen in the future. Sometimes that fear is rational and sometimes not. And sometimes it's about something that will happen in three minutes (stepping onto a stage to make a presentation, for example) or in 30 years (having enough money to retire).

In the United States, anxiety is the most common mental illness, affecting more than 40 million adults each year. Data from the National Institute of Mental Health has indicated that about 30% of Americans experience clinical anxiety at some point in their lives. Globally, according to the Institute for Health Metrics and Evaluation, an estimated 284 million people had an anxiety disorder as of 2017, making it the most prevalent mental disorder worldwide. And recent workplace data from Mind Share Partners, SAP, and Qualtrics suggests it's widespread on the job: Nearly 37% of workplace respondents reported symptoms of anxiety in the past year. These numbers will only increase in the wake of the pandemic.

Idea in Brief

Anxiety isn't always counterproductive. It can prompt you to react quickly to threats, and in a crisis it can spur you to make your team more resourceful and productive. But unchecked, it zaps your energy and clouds your decisions. Anxiety is a powerful enemy; you need to disarm it. Morra Aarons-Mele offers advice on how you can inspire and encourage your team even when you're struggling yourself:

- First you must accept your anxiety and know what triggers it. Once you do that, you can begin managing it.

- Work to distinguish between the possible and the probable.
- Structure your time, take small, meaningful actions, and try techniques that reduce anxiety's physical symptoms.
- Set yourself up so that anxiety doesn't harm your ability to lead.

The good news for those of us who have managed anxiety for a long time is that we were made for this moment. Data shows that anxious people process threats differently, using regions of the brain responsible for action. We react quickly in the face of danger. We may also be more comfortable with uncomfortable feelings. When channeled thoughtfully, anxiety can motivate us to make our teams more resourceful, productive, and creative. It can break down barriers and create new bonds.

So anxiety isn't useless. In an economic crisis, the anxiety that keeps us up at night may help us fathom a solution to keeping our businesses open. But left unchecked, anxiety distracts us, zaps our energy, and drives us to make poor decisions. Anxiety is a powerful enemy, so we must make it our partner.

Whether you have a diagnosed anxiety disorder or are having your first dance with this intense emotion, you can still be an effective leader. But I'll be blunt: If you don't look your anxiety in the face at some point, it will take you down. This isn't easy, but doing it will change your life and your ability to lead others for the better.

So today, in this especially anxious moment, let's begin. The first stage is learning to identify your anxiety: how it manifests itself and how it feels. The second stage is taking action to manage it both day-to-day and in challenging moments. The third stage entails making smart decisions and leading others in anxious times. Finally, the fourth stage involves building a support infrastructure to help you manage your anxiety over the long term.

Acknowledging and Accepting Your Emotions

A common coping mechanism for leaders is to push through stress, fatigue, and fear. But that's succeeding *in spite of* your emotions, when it's far better to thrive *because of* your emotions. You have to learn to accept

your anxiety—even though this may seem uncomfortable or counterintuitive.

Label what you're feeling

Angela Neal-Barnett, an award-winning psychologist, expert on anxiety among African Americans, and author of *Soothe Your Nerves*, is a firm believer in being honest with yourself. When you name a feeling—by saying to yourself "I'm anxious"—you can begin to address it. You can learn how anxiety informs your behavior and your decisions and what causes it to surge, which will equip you to manage it.

No one has to hear you say it. This is for you. Take the time to wallow in your thoughts. Let yourself experience the discomfort of fear and anxiety. Play out worst-case scenarios in your head. Allow your imagination to go wild with catastrophe. Cry. Grieve. But don't turn away. As Alice Boyes, a former clinical psychologist and author of *The Anxiety Toolkit*, says: The more you try to control your anxiety, the more it fights back.

Decades of research on emotional intelligence have shown that people who understand their own feelings have higher job satisfaction, stronger job performance,

and better relationships; are more innovative; and can synthesize diverse opinions and lessen conflict. And all those things make people better leaders.

If the word “anxiety” feels wrong to you, label it whatever you like. Call it “unease” or “temporary uncertainty” or even give it a silly name. I think of my own anxiety as a separate character who travels with me. She doesn’t have a name or a face, but I know when she’s present.

The leadership coach and CEO of Reboot, Jerry Colonna, says that the best way to deal with uncomfortable feelings is to welcome them in. Think of your thoughts and emotions as trains coming in and out of a station, he advises. Watch them arrive and depart without attachment. Imagine saying, “Hello, anxiety. See you later, fear.” This technique actually will help you build distance from the negative feelings in your mind.

Sometimes it may be impossible to get rid of your anxiety, which can feel frustrating. Rebecca Harley, a psychologist at Massachusetts General Hospital and Harvard Medical School, emphasizes, “The goal is not to magically make things perfect. The goal is to learn to surf

the waves of distress successfully. Give yourself credit even if things don't feel all the way better."

Play detective

Once you've labeled your anxiety, you can start pinpointing when it appears and why. Harley helped me learn to do this. When you feel anxious, take note of your physical reactions—what she calls the "early warning system" that anxiety might be taking over.

Your triggers might be small. You might notice a stomach flip and a spark of dread when you see someone's name pop up in your in-box. Or they might be bigger. When unemployment numbers skyrocket, you might feel nauseous and unable to focus even though you still have a job.

When an interaction or a situation sets you off, examine why. You might be hesitant to delve into issues from your childhood, but "unresolved business" from your past, as Colonna puts it, is very much present in—and relevant to—how you lead. He notes it can be a relief to truly understand how your old wounds inform your present behavior. When I realized that my near-constant worry about going broke stemmed more from my childhood

than from my current financial situation, I was finally able to proactively manage my money, after years of avoiding it and piling up debt. I broke a damaging pattern.

It's also good to understand how you react when triggered. I call these anxiety "tells." Social worker and therapist Carolyn Glass suggests asking yourself, "How did I respond to that anxiety in that moment? And were those behaviors helpful or not? Did those behaviors fuel or alleviate my anxiety?" Glass says that writing down your fears will help you examine them. Keeping a journal of your anxiety—when it happens, what triggers it, and how you reacted—is a great way to develop self-awareness. Your tells may not always be negative behaviors, though; for instance, many of us find ourselves connecting with friends and family more during stressful times. When I'm very anxious, I cook and freeze meals!

Many successful leaders react to anxiety by working harder, holding themselves and others to an impossibly high standard, or trying to control things that are beyond their power. For them, it's hard to imagine *not* fussing over every project and detail in their work lives, not taking responsibility for everything or always giving their

all. “People respond to anxiety by trying to be more perfect and more in control,” Boyes says. “They not only have a Plan B but Plans C, D, and E.” In many societies those behaviors are rewarded. We think of it as a “good work ethic,” but often perfectionism and overwork only cause further anxiety—in yourself and others.

Imagine a CEO who is terrified by the economic news surrounding Covid-19. He jumps into the problem in the way that’s worked for him in the past: making detailed projections on all aspects of the business. He buries himself in these charts while constantly consuming news about the crisis. Some of his team might wonder what he’s up to or feel unsettled by his visible yet unspoken panic. Are the charts he furiously creates accurate? Who knows?! But the deep dive into worst-case scenario planning gives him the illusion of control.

Your tells may also be physical. Anxiety can manifest itself as tightness in the chest, shallow breathing, clenched jaw muscles, frozen shoulders, gastrointestinal symptoms, skin breakouts, appetite changes, and radical shifts in energy. When I recently had a panic attack, for example, I was convinced it was heart failure—even though I’d had panic attacks before.

To help you identify the ways anxiety may be physically affecting you, try this two-part exercise:

First, sit upright in a chair. Put your feet flat on the floor, and your hands on your lap. Keep your chin neutral. Note which part of the body you can immediately feel. Then, with your eyes closed, scan through the following:

- Your head
- Your jaw
- Your neck cords
- Your shoulders
- Your wrists and forearms
- Your chest
- Your upper back
- Your lower back
- Your stomach
- Your hips
- Your hamstrings and rear
- Your calves, ankles, and feet

Note which ones feel tight, and to gain some relief, breathe into the areas of tightness or pain.

You can also pay attention to what's happening with your body at different points during the workday, when specific events occur, or when you make certain decisions:

- How do you feel at 9 a.m., noon, and 3 and 6 p.m.? Does your body change over the course of the day?
- If you get stressed, does a particular part of your body react?
- How often do you rely on a drink, drug, muscle relaxant, or over-the-counter pain relief over the week?
- Does your body feel different after you exercise? Do your shoulders feel lighter?
- How does your body feel on the weekend or when you're doing something you enjoy?

Sort out the probable from the possible

Once you understand your triggers and tells, you can start developing a new relationship with your anxiety.

Remember, some anxiety is rational and helpful. In an economic downturn it makes sense for a leader to feel anxious. You might have to lay people off. Your business

might fail. But you might find that you get stuck in a negative thought loop that prevents you from moving forward; you start obsessing. Boyes points out that some leaders get so focused on the worst-case scenario and overwhelmed by scary possibilities that they become frozen.

So how do you avoid being stuck? Here I turn to advice from Colonna: “Differentiate what’s possible from what’s probable. It is *possible* that everyone I love will die of a pandemic and I will lose everything I hold dear. But it’s not *probable* that everything that we love and hold dear will disappear.” Try to distinguish your worst fears from what is likely to happen. This will help calm you and give you space to move forward. So when a catastrophic thought comes into your head, such as “My partner and I are both going to lose our jobs” or “I’m definitely going to get sick,” remember that you’re an unreliable narrator when you’re anxious. Check in with someone else you trust and ask for that person’s help in telling what is likely to unfold from what is a long shot.

Back in early March, when the stock market first slumped and people’s fears about Covid-19 spiked, one of my biggest clients canceled work with my small business.

I quickly convinced myself that our company was doomed, that it would be only a matter of months before we had to close up shop. “We’ll never survive this,” I kept telling myself. But then I consulted my business partner—a more reliable narrator than I—and she suggested we readjust our forecast, which we did. Now we’re projecting that we’ll lose half our revenue for the year. This is probable and upsetting, but it’s far different from going out of business completely.

Focusing on what’s probable also takes flexibility—the future won’t be what you thought, and that hurts. When my preschooler really wants to keep coloring, but it’s time for dinner, I ask her, “Please be flexible. You can color later, I promise.” I’m now trying to do what I’ve taught my kids for years: to handle the disappointment of things not going the way I expected or wanted. These disappointments are real, and sometimes the changes are grave. Acknowledge the grief and anger you feel (at least to yourself) and then make adjustments, identifying the aspects of your vision that may still work, and focus on what’s probable.

Taking Action to Manage Your Anxiety

Once you make your way through these three steps, you can start to manage your anxiety daily in ways that allow you to grow as a leader and be more resourceful and productive.

The following tactics can help ground you.

Control what you can

Many faith traditions teach us to accept what we cannot control, without preoccupation or panic. But in the middle of an anxiety attack at work, you probably don't have time for philosophy. So here's what to do when things feel completely off the rails.

Structure your time. A solid body of research shows that improved "time management disposition"—meaning your attitude toward how you organize and value your time—has a positive impact on mental health. And it's especially crucial when you're gripped by anxiety.

First thing in the morning, create a to-do list and a detailed schedule for your day. I like to do it while having my coffee. You might use 30-minute increments to spell out when you'll shower, take a lunch break, make a phone call, or tackle that report that needs to get done. This is

what many experts call “timeboxing.” While you’re at it, try to avoid what cognitive behavioral therapy terms “cognitive distortions.” These are the catastrophic thoughts, self-judgments, and all-or-nothing ideas that often accompany anxiety.

Be careful not to overschedule or overestimate your productivity; instead focus on the critical work and leave time to take care of yourself.

Take small, meaningful actions. During the first few weeks of the coronavirus shutdown, traffic dropped drastically where I live. The local department of public works took that time to repaint all the crosswalks. For a week, roads were halfway blocked off as DPW crews painted. It wasn’t a big deal because our normally bustling town was quiet. And each time I slowed down to drive past one of the crews, I smiled because it struck me: This is their small, meaningful action.

When you feel anxious, an immediate task can easily become overwhelming. Take running a cash flow analysis for your business. When you open up the accounting software, your mind might go to a dark place, and all of a sudden a month’s worth of figures have spiraled into the

business tanking and your losing your home. To break that mental spiral, take a small, meaningful action. If running a cash flow projection terrifies you, organize some receipts or clean up some file folders until the panic subsides.

In general, focus on the near term whenever you can. You may not be able to tell your employees what will happen next year—or even three months from now. You can't promise everything will be OK. But you *can* help your people be safe this week. Focus on that, and then deal with the big questions when you feel calmer or when you can get input from trusted colleagues. Sometimes you have to turn off the future for a little while and just manage through the present.

Develop techniques for situations you can't control

Of course, it's not always possible to turn off the future. What if your board needs those cash flow projections in the next 30 minutes and you're in a downward spiral? Here you'll want to have tools that help you calm down quickly so that you can get your job done.

Find a mindfulness technique that eases your acute anxiety

Neurologist Victor Frankl famously said, “Between stimulus and response there is a space. In that space is our power to choose our response. In our response lies our growth and our freedom.” This is mindfulness in a nutshell. Even if you are high on anxiety and short on time, you can claim the space in between.

There are lots of ways to do this; the key is to find what’s most effective for you. One option is to focus on your breathing. Belly breaths are a classic technique. Others prefer what’s called “the 4-7-8 method.” Either is simple to memorize and subtle enough to do at your desk. When you deliberately slow your breath, it sends a message to your brain to calm down, and your brain then sends the message to your body so that many of the physical symptoms of anxiety—such as increased heart rate and higher blood pressure—decrease.

You can also shift your attention. Glass says this technique is “great for someone who doesn’t want to meditate but gets maladaptively anxious and cannot focus on anything else.” Focus first on your anxiety, and then slowly turn your attention to something tangible,

something you hold in your hand, like a book. By concentrating on an object in the present moment, you can turn the volume of your worry down until it's background noise.

If I'm full of anxious energy and unable to sit still, or if quiet breathing exercises don't work, I like to loudly blast a favorite song and dance for five minutes. Some people like to sing instead. Experiment with what works for you and then keep that tactic in your back pocket for when you need it.

Compartmentalize or postpone your worry. Sometimes I talk out loud to my anxiety, saying, "Sorry, I'm going to deal with you after I finish my work." You may want to write the worry down and save it for a specific time—maybe later that day or your next session with your therapist.

In times of crisis you may actually find that things that worried you in the past fade into the background. The urgency of what's happening in the moment takes over. To stop your anxiety from sneaking into the foreground, you might tell it, "You can stay where you are. I'm part of the solution here, and I need to get this task done."

Make a connection. Connecting with others can break the negative thought loop that often accompanies anxiety. Instead of focusing on yourself, you turn your attention outward. When I asked my friend and colleague Cheryl Contee, the CEO and cofounder of the digital agency Do Big Things, how she was staying motivated during the crisis, she said that she was trying to "be a good neighbor," something she learned from her grandfather, William G. Contee, who has a park dedicated to him in his Baltimore neighborhood. "Being a good neighbor is surprisingly simple—it's just about connecting on the human level," she told me. "Do you say hello to your neighbors? Have you asked how they're doing or if they need anything?"

Contee also connects digitally with people in her field, who support one another and contribute to causes they care about. At her company she and her colleagues are leaning into talking about their feelings and families, doing a lot of checking about how to balance homeschooling with work. "We're all veteran virtual knowledge workers, but having kids around and being responsible for their education is a new challenge we're facing together," she said.

In your own life, think about performing a quick, generous act. You might check in on a former colleague via text message. Or ask a family member how you might help. When I'm feeling anxious, I sometimes go to LinkedIn and "like" articles written by my colleagues or write up an endorsement of their work. This gets me out of my head and focused on something more positive.

Finally, if anxiety is persistent and hampering your days, you might consider consulting a therapist or mental health professional. Talking to someone trained in helping others manage anxiety may give you additional coping mechanisms to address debilitating symptoms.

Limiting Anxiety's Impact on Your Leadership

Once you have a better sense of how you experience anxiety and how you can manage it daily, it's time to turn to how it affects your leadership and management abilities.

Make good decisions

Anxiety can impair our judgment. It can cause us to focus on the wrong things, distort the facts, or rush to conclusions. Ideally, we could postpone critical decisions until we're in a better frame of mind, but that's not always possible.

In anxious times it's important to proactively set yourself up to make good choices. Much as you do when separating the possible from the probable, start by acknowledging that your emotions can make you an unreliable narrator and that you will likely be prone to negative thoughts. Let's say you're prepping for a speech and the last time you spoke to a group of a similar size, you felt that you bombed. You may even have a long-held belief that you're a terrible public speaker because a middle school recitation drew snickers. Ask yourself: Are you being objective? If you're not sure, check whether your memory is correct, perhaps by asking a colleague who was in the room for feedback.

Of course, you need to ask the right people. Boyes suggests you find a trusted adviser with a decision-making style that differs from your own. If you're impulsive, consult someone who is methodical and conservative, for example.

Ultimately, every leader should develop a team of “real talk” peers: people who will provide their unvarnished opinions. You can fill this role for others, too. You can still offer them clarity and insight even if you’re an unreliable narrator of your own experience.

Practice healthy communication

One of the most dangerous aspects of anxiety is that it’s contagious, and leaders set the tone. Daniel Goleman, the renowned psychologist and author of *Emotional Intelligence*, calls this “neural Wi-Fi,” in which humans pick up on others’ unspoken feelings.

If you’re not admitting that you’re anxious but instead emitting irritability or distraction, you’re not doing your staff any favors. But how can you be honest with your people in a way that doesn’t strike fear in them? What degree of emotion is appropriate to express?

Ultimately, how much you disclose is a personal decision. As an owner of a business and the host of a podcast about anxiety and mental health, I tend to be an open book. But I know that most leaders don’t share their demons. Few feel comfortable starting a staff meeting with “Wow, I’m anxious today.”

But self-aware leaders know when it's appropriate to be vulnerable. And here's the thing: Your staff needs you to be transparent and honest about anxiety and mental health, especially when the future of your company and their livelihoods are uncertain.

Amelia Ransom, the senior director of engagement and diversity at Avalara, says that she wants her leaders to admit when they're not doing OK, because it affirms her experience. "It makes me feel normal if someone I respect and trust admits they aren't all right. I think, 'Thank you for being human,' and I want to follow that person." Ransom recounts a powerful moment when a senior executive in her company brought the staff together on a videoconference and said, "I can't tell you, 'You got this.' What I can do is hold space for us to be together right now, to talk and figure some things out."

Admitting "I'm anxious today" or "I didn't sleep well" lets everyone else in the room breathe a little easier. ("Phew, it's not my fault he is so tense.") And remember, you don't have to share details; just share the state you're in.

The social psychologist Amy Cuddy tells us we need leaders who exhibit both warmth and strength. "Most

leaders today tend to emphasize their strength, competence, and credentials in the workplace, but that is exactly the wrong approach," she writes. "Leaders who project strength before establishing trust run the risk of eliciting fear, and along with it a host of dysfunctional behaviors." Nothing establishes trust more effectively than the emotional connection fostered through empathy and shared humanity. This is why being open about your own anxiety can be so powerful. It builds trust when you can ask teammates, "How are you?" and they don't feel as if they have to lie or put on a happy face, because they know you feel the strain, too.

This doesn't mean that you fall into a puddle of tears during a videoconference, of course, or visibly lose control. And while your workers might want to know that you're closely monitoring cash flow to make sure bills get paid, they don't need to know that your anxiety is deeply rooted in your parents' money troubles during your childhood. It's possible to model taking care of your mental health without making people lose confidence in your competence.

Imagine you're in an anxiety spiral from reading news about Covid-19, but you need to lead a staff meeting in 10

minutes. You could open the meeting by saying, "Obviously, the news is getting more upsetting by the minute, but I want us to put that aside for the next half hour while we go through this call." Or you could be even more vulnerable and share that you're working to contain your scary thoughts by giving yourself what Glass calls a "worry hour," when you allow yourself to indulge your biggest concerns before putting them away again and forging on.

If you want to encourage people to share but don't want the conversation to slip into an anxiety fest, you can use a red-yellow-green exercise. Team members individually indicate where their moods are that day with one of the three colors, and they can expand on why if they wish. This allows people to share if they feel comfortable doing so and gives you useful information about the emotions of the group. You can then adjust your communication style and messaging accordingly.

And remember, while being positive is important to prevent emotional contagion, you don't want to give anyone false hope. If you get tough questions like "Is my job safe?" or "Will we be in business in six months?" it's not your job to divine the future. No one has a crystal ball,

and so you can say what you know to be true in this moment and affirm the importance of working together and focusing on what each person can control.

Building a Support System

The final step in leading through anxiety is making sure you have ongoing support. This means not only surrounding yourself with the right people but also developing routines that help you deal with bouts of anxiety and lay the groundwork for maintaining your mental health.

Schedule, structure, and scenario plan

When you have anxiety, you need to be intentional about what your days look like, as I discussed earlier. The methods are basic: making lists, prioritizing, and breaking work into manageable chunks. Chop tasks that make you extremely anxious into bearable pieces. I learned this trick from my own psychiatrist, Carol Birnbaum.

Also use the detective work you did about your triggers to prepare for situations or events you know will cause you anxiety. If public speaking stresses you out, make

sure you leave plenty of time to rehearse presentations. If you're afraid of flying, mentally rehearse a business trip from "I'm going to pack" to "I'm going to order a cab and call my friend while I'm on my way to the airport" to "I'll buy M&M's when I get there because they make me happy." And finally, once on the plane: "I'm going to take a Xanax, do a calming meditation, and survive."

I get anxious when I'm working far from home and haven't heard from my nanny or husband. I worry something bad has happened and get distracted from what I'm supposed to be doing. To counter this I ask my husband or the babysitter to text me with an update every three hours. That way I don't pester them when they might be driving with the kids in the car, for example. And knowing that they will keep me updated allows me to sink into my work.

Know who your "safe team" is

Since you want to spare your employees the messy details of your anxiety, you need a place for those emotions to go. Make sure you have a "safe team" of people to whom you can confess scary thoughts. They can include a therapist, a coach, a mentor, a spouse or partner, and

friends. It could be an intimate group of fellow leaders, online or off-line, who commit to sharing in confidence and making space for one another's difficult emotions.

Practice self-care

I don't need to belabor this point. You know what self-care means for you, whether it's sleep, exercise, hobbies, massage, spending time alone, or being with people you love. The point is, take it seriously, as if your doctor had written you a prescription for it. It's neither frivolous nor optional for you as a leader. And aspects of it you feel comfortable sharing can benefit your team: When you model good practices, others feel permission to take care of themselves, too. This could be as simple as letting people know that you don't take your phone upstairs when you head to bed, that you're taking an hour during the workday to exercise, or that you're limiting exposure to news or Twitter.

Putting in place the support infrastructure to manage your anxiety will help you ride out setbacks and tough times. It's a strategy for long-term success and sustainability as a leader. It means you'll have better

workdays, both when things are status quo and during transitions and tough times.

Ultimately, anxiety comes with the job of being a leader. The process of managing it can make you stronger, more empathetic, and more effective. It just might be bumpy along the way. So remember to treat yourself with compassion. Recognize that you're doing the best you can, that your emotions are normal, and that the healthiest thing you can do is to allow yourself to experience them.

Far too many of us think it's taboo to talk about mental health at work. I know many leaders who don't feel as if they can walk into a staff meeting and say, "I'm anxious today."

Why not? And why not now? These are not normal times, and acknowledging a universal emotion can help people understand that what they're feeling is OK.

We're in desperate need of better models of leadership, especially when society tells us that anxiety and depression are weaknesses. The data bears this out: A 2019 Mind Share Partners report found that 86% of U.S.

job seekers thought it was important for an employer's culture to support mental health, but only 37% of employees saw their company leaders as advocates for mental health at work.

This time of crisis—in which those of us with a history of anxiety may be experiencing it acutely while others may be feeling it intensely for the first time—is an opportunity to change that perception.

You can play a role in telling a different story.

When Anxiety Becomes Unbearable

by Gretchen Gavett

Maybe you've been anxious all your life. Maybe the Covid-19 pandemic has taken it to an unprecedented level. Or maybe you don't suffer from anxiety but work with people who do.

Even if you're used to managing it on a day-to-day basis, it's not always clear when anxiety is cause for

concern. So I asked **Dr. Ellen Hendriksen,** a clinical psychologist and the author of *How to Be Yourself: Quiet Your Inner Critic and Rise Above Social Anxiety*, about the warning signs that someone is really struggling.

Dr. Hendriksen offered advice on addressing heightened anxiety with humanity and compassion and, when the person struggling is a work colleague, discussing it without overstepping. Our conversation, which took place over email, has been edited for clarity.

HBR: ***What are the signs that your anxiety is reaching an unsustainable level?***

Hendriksen: We know anxiety has escaped its confines and is running wild when it surpasses the thresholds of distress or impairment. *Distress* means intense stress is overwhelming your usual ways of coping. Maybe you've always been able to manage your anxiety with yoga, a sense of humor, or some healthy perspective, but now nothing seems to keep the lid on. *Impairment* means the anxiety is getting in the way of living your life. For instance, you can't focus, so you're behind on your work, are losing sleep, or are so preoccupied that you can't be present with your kids or partner.

What should you look out for in your boss, colleagues, or employees?

Problematic anxiety is mostly internal and therefore is hard to spot. However, clues include unmanageable worry or irritability, inability to focus or concentrate, and physical restlessness (pacing, being on edge). In conversation, or on a Zoom call, you can often sense the person has tunnel vision—they might be hyperfocused on a worry, keep coming back to a topic, or refuse to consider others' points of view.

In addition, problematic anxiety often leads to under-control or overcontrol. People who under-control can be passive or all over the place. Their actions can be scattered, inefficient, and unhelpful—you might find yourself thinking, "What are they *doing*?" They either follow every impulse or give up altogether, bending as the wind blows.

Overcontrol can take the form of micromanagement, rigidity, hypervigilance for potential threats, refusing to try new approaches or adapt, or insisting that there is a right way to do things. There are the folks who, during a time of high anxiety, throw themselves into work or

anything else they can control, from financial planning to making a spreadsheet of their canned goods.

But remember: Overcontrol is a problem only when it causes distress or impairment to the anxious person or to those around them. Focusing on work is fine if it gets someone's mind off the crisis, but if work becomes their *only* focus, and especially if their health or relationships suffer as a result, they have crossed the line into unhealthy overcontrol. Keep this in mind for yourself as well.

Pretty much everyone has a relatively high level of anxiety at the moment. How can you differentiate that from something more troubling?

It is important that someone's response isn't disproportionate to the threat. Before Covid-19, it would have been over the top to wear gloves and a mask while grocery shopping. These days everything has changed. A higher-than-normal level of anxiety is expected and appropriate. But the disproportion rule still applies: For example, refusing to go to the emergency room if your appendix bursts, or hitting the grocery store in full scuba gear, is still generally considered to be over the line.

Give everyone plenty of slack, understanding, and validation. Also, when someone seems especially anxious, keep in mind that you may not know the full context. Maybe their sister is a nurse on the front lines in New York. Maybe their spouse was laid off and they're worried about paying the mortgage. Maybe they have an underlying condition that puts them in a high-risk group. Their anxiety may not be disproportionate at all.

What are some things you shouldn't say to a colleague who seems anxious? And what can you do that might actually help?

Don't try to offer quick fixes: "Have you tried yoga?" or "I hear lavender essential oil can work wonders." Advice like that, while well-intentioned, comes across as invalidating—"Oh, is that all I have to do? Silly me!" Worse, offering advice creates an expert/amateur dynamic, rather than a relationship of equals.

Likewise, dismissals like "Calm down," "There's nothing to be afraid of," or "Just don't worry about it" feel invalidating and unsupportive.

Many people feel uneasy about offering help to a colleague they're not close to. But no matter the depth of

your relationship, you can always validate their experience ("It totally makes sense that we're all stressed right now" or "Trust me, nobody is doing their best work these days") or make a workplace-appropriate disclosure of your own ("It's been a real challenge to juggle everything" or "The worst part for me is not knowing how or when all this will end").

Some worries are expressed as "what ifs"—"What if I get quarantined and can't work?" or "What if my elderly parents get sick?" Know that the "what if" is rhetorical, but go ahead and inquire about an answer: "That's a scary thought. What *would* you do?" Anxiety is driven by uncertainty, and generating a plan creates certainty, which in turn can reduce anxiety. Supporting your colleague as they think through a plan of action (without proffering advice—refrain from "I'll tell you what worked for my brother") can be helpful without invalidating their fears.

What should you do if you, a colleague, or an employee has a panic attack?

Panic attacks feel awful. It sounds silly, but remind yourself (or your colleague) that it's a panic attack. It's

easy to get caught up in the feeling that you might be dying, are having a heart attack, or have finally snapped and lost it. But remember, it's a panic attack, and panic attacks always end. What goes up must come down.

Next, if you or your colleague has prescribed medication for panic, this is the time to take it.

If you're at home and have a panic attack, a nonmedication response is to fill a sink or a big bowl with cold water, add ice cubes if you can, and dunk your face in. Hold your breath and keep your face immersed for 30 seconds. This triggers the *diving reflex,* which is an evolutionary response that shuts down all nonessential body functions—including strong emotion—during a fall into cold water. It kicks in the parasympathetic nervous system and calms you down. Alternatives include taking a cold shower or putting an ice pack over your eyes and holding your breath for 30 seconds.

If you're trying to help a colleague, don't chatter anxiously at them or pepper them with questions. Stay as calm as you can and ask them to breathe deeply and—most important—slowly! Rapid breathing can mimic hyperventilation and make a panic attack feel worse. Ask them to inhale for a count of six and exhale for a count of

10. This takes advantage of a natural physiological phenomenon called *respiratory sinus arrhythmia,* which means your heart beats faster on an inhale and slower on an exhale. If you exhale for longer than you inhale, over time your heart rate will slow, which in turn will calm your other body systems.

What if someone's workplace performance suffers over time because of continued anxiety? How should managers and employers respond?

Managers can't and shouldn't ask about private health information. But they can address specific tasks or behaviors. If deadlines are being missed or projects are falling perilously behind, you can start a conversation. Use the same tone you would in asking about a physical illness or injury—be caring and open. Tiptoeing makes things awkward and can backfire. Say: "I wanted to check in given the last few deadlines. I know this crisis has been a big challenge for everyone. You don't have to navigate this all by yourself. Let's talk about how we can support you." Or say: "I want to acknowledge that things have been really stressful and uncertain lately. It's not your style to let things fall behind, so I wanted to check in.

You're such a vital and extraordinary part of this team; let's make sure you have what you need."

This is a great time for both you and your employees to take advantage of any mental health benefits your workplace offers, like an employee assistance or behavioral health program. During this crisis, many therapists are offering telehealth sessions through health care–compliant platforms. And even though privacy can be hard to come by with kids and partners within earshot, enterprising folks have taken their laptops out to their parked cars for online therapy sessions, or talked to their therapists on the phone in our nearly empty streets.

What happens now that many of us are working from home? How can you stay on the pulse of your employees' anxiety virtually, without being too invasive or violating anyone's privacy?

It's OK to be more direct than usual given the circumstances. Working remotely, it's harder to pick up on the same signs you would see in person.

Be transparent—acknowledge that it's a difficult time for everyone and you want to check in. Ask how they're doing and how you can help. If they say, "I'm fine," don't

respond with “Great!” and mentally check off the box. Instead, respond with a spirit of flexibility and openness: “If that changes, let me know” or “I’m so glad—let’s keep checking in as the weeks go by.” Everything that makes you a good human—sincerity, flexibility, caring about your colleagues—will make you a good manager in this unprecedented time.

---

5 Ways Leaders Accidentally Stress Out Their Employees

by Tomas Chamorro-Premuzic

Decades of scientific research show that stress and anxiety are prevalent problems at work, contributing to deficits in employee morale, well-being, and productivity. While anxiety is caused by a range of factors, including issues unrelated to people’s jobs, one common and pervasive cause *is* something specific to the workplace: incompetent leadership.

Managers and leaders have a direct effect on their employees' stress and anxiety levels. What they say, feel, and do hugely influences their team's physical and emotional well-being. And the more senior leaders are, the more people they are likely to influence—positively and negatively.

But sadly, far too few leaders are aware that they have this power. And many are overconfident in their leadership skills, creating a gap between their perceived and actual levels of competence. This explains why even well-meaning bosses may inadvertently contribute to high anxiety levels in their employees and have a limited capacity to correct and improve their behavior: If you think you are leading effectively, what is the point of changing?

It is for this reason that leaders must pay a great deal of attention to how they act and communicate. The importance of this is exacerbated during times of increased uncertainty, as we often look to leaders to guide us in the face of fear, to provide us with clarity and direction, and, most of all, to give us reasons to remain hopeful and optimistic.

If you are a manager or a leader, it is useful to internalize some key psychological lessons about how your behavior—what you say, do, feel, and express—impacts your team, especially when you are not aware of it. In particular, there are five behavioral patterns that most often increase people's anxiety level. If you can spot them, you can learn how to change them in order to become a more effective leader.

1. The use of negative language

Too often we focus on nonverbal communication as a signal for conveying emotions—how we move our hands or which facial expressions we make—when in reality, the words we say are more likely to convey how we feel and what we think. As the growing field of algorithmic text mining and natural language processing shows, there is a systematic and robust connection between the type and frequency of words we choose to express ourselves and our moods and temperaments.

This means that even when you think you're discussing your business strategy dispassionately, the way you talk about it and the language you choose will convey your emotional and mental state to others—irrespective of

your intentions. Leaders in particular can expect the emotional impact of their words to be even stronger when they are written. People tend to reread important messages, internalizing their affective content.

Research has shown that to avoid accidentally triggering anxiety through language, best practice is to refrain from using *negative words* (for example, *horrific, shocking,* and *dangerous,* as well as euphemisms such as *challenging, problematic,* and *undesirable*). In fact, the only criterion for determining whether a word is negative is whether it increases the listener's negative affect—in other words, that it might elevate their levels of anxiety, worry, and concern. Even if two leaders are in the same situation and describing the same state of affairs, they will have a different effect on the public if they talk about "hope," "improvements," or "light at the end of the tunnel" as opposed to "death toll," "mortality rate," or "depression."

2. Unusual or erratic actions

We often celebrate spontaneity and unpredictability as critical ingredients of creativity, as if they were integral components of a free spirit. In reality, however, most

people want to eliminate as much uncertainty and unpredictability from their lives as they can, as both tend to trigger anxiety.

The Covid-19 pandemic makes this clear. We are shocked not only by the virus's devasting effects on our lives but also by our inability to predict what will happen. There is not much certainty leaders can provide when they are equally unable to predict the future. But they can, at the very least, avoid being an additional stress agent by acting in consistent and predictable ways.

If you are a boss, don't introduce an unnecessary layer of complexity to your employees' lives by making them guess what you will do next. Be reliable, predictable, and even boring if necessary. You may be the only predictable factor your employees can count on in a time of great uncertainty.

In simple terms, this means providing a clear structure to your meetings and communications, sharing expectations up front, avoiding last-minute changes and cancellations, and, wherever possible, continuing with the same routine you had before the crisis or big change.

3. Emotional volatility

Excitable bosses are like a roller coaster—they may be fun for sensation seekers, but they are stressful for almost everyone else. The last thing your employees want during difficult times is to see emotional volatility in their leaders. It is a bit like provoking someone into an emotional discussion when they've had a really bad day—it is not going to end well.

This may be easier said than done, but being a leader requires a certain level of competence for dealing with pressure. Especially in a crisis, remember that your own stress will only amplify other people's stress. The main implication here is that you should work very hard to manage your impressions, contain your emotions, and put on your best poker face in front of your employees.

What might this look like? In general—and this goes back to the second point—your team is looking to you for stability and guidance amid the chaos. If you are typically calm and stable, try to remain so as much as possible. Even if it may be normal to display some degree of emotional volatility during a crisis, the fewer changes your team perceives from your typical patterns of behavior, the less stressed they will be. If your natural style is volatile and reactive, however, you may be better

off projecting an aura of calmness and composure, as if you had just taken up meditation. This shift may feel extreme to you personally, but over time it will help you better tame or filter your own anxiety. Once your team begins to notice the change, they may feel less on edge themselves, too.

Actions that have been found to mitigate emotional volatility include a regular practice of mindfulness, frequent exercise, better sleep quality, and internalizing feedback from others so that you realize when you may be derailing.

4. Excessive pessimism

We live in a world—especially in the West—that stigmatizes negativity and condemns pessimism as if it were a psychological problem. In fact, pessimism is underrated, as it helps leaders to detect and prevent potential threats, minimize risks, and avoid arrogant and overconfident decisions. That said, during stressful and anxious times leaders' pessimism is more likely to turn into a liability, demotivating others and pushing their already high anxiety to stressful levels.

This is why, even when you cannot find reasons to project optimism, you should still refrain from displaying outright pessimism. Even if your natural response is to feel pessimistic, projecting this onto others may further their anxiety. Being able to control it and project calmness and composure will strengthen your colleagues. Remember that leadership is not about you; it's a resource you provide to help others.

Further, because it's likely that your team expects a certain degree of optimism from you during uncertain times, they may already discount for this. If you say things will be great, they will believe you; but if you tell them things will be bad, they may interpret the situation as worse than it is.

5. Ignoring people's emotions

Perhaps the biggest mistake you can make during stressful times is ignoring your team's emotions. This error often occurs when a leader is hyperfocused on dealing with their *own* emotions. While you need to understand your own anxiety and get it under control, it is also critical to manage how others are perceiving your well-being. If they think you cannot manage yourself,

they won't trust you to manage them. The key here is empathy: You will only succeed if you are focused on the people around you, not on yourself.

In the past two decades, a great deal of research has highlighted the key role that emotional intelligence (EQ) plays in developing empathy. More specifically, we have learned that leaders with a high EQ are better at understanding and influencing other people's emotions, as well as controlling their own. Some leaders are naturally better at this than others. Unfortunately, no one will suddenly wake up with a higher EQ overnight. But they *can* work on their willingness to understand other people.

A critical starting point is remembering that during difficult times it is more important to monitor people's affect, mood, and stress rather than check on their work performance, productivity, or task management. Simple ways to achieve this are to have more one-on-one meetings with team members, increase the frequency of your communication, ask open-ended questions that invite people to engage, and show empathy whenever possible. As the great Dale Carnegie put it, "When dealing

with people, remember you are not dealing with creatures of logic, but creatures of emotion."

In short, you will be less likely to increase anxiety in others if you make a commitment to thinking more deeply about how your actions impact them. As a leader, you are an amplifier of people's emotions. If you do things right, you can bring out the best in people even in the worst of times. If you do things wrong, you will lower morale and performance even when things are fine.

Originally published in May 2020. Reprint H05LK7

When Machine Learning Goes Off the Rails

by Boris Babic, I. Glenn Cohen, Theodoros Evgeniou, and Sara Gerke

WHAT HAPPENS WHEN machine learning—computer programs that absorb new information and then change how they make decisions—leads to investment losses, biased hiring or lending, or car accidents? Should businesses allow their smart products and services to autonomously evolve, or should they "lock" their algorithms and periodically update them? If firms choose to do the latter, when and how often should those updates happen? And how should companies evaluate and mitigate the risks posed by those and other choices?

Across the business world, as machine-learning-based artificial intelligence permeates more and more offerings

and processes, executives and boards must be prepared to answer such questions. In this article, which draws on our work in health care law, ethics, regulation, and machine learning, we introduce key concepts for understanding and managing the potential downside of this advanced technology.

What Makes Machine Learning Risky

The big difference between machine learning and the digital technologies that preceded it is the ability to independently make increasingly complex decisions—such as which financial products to trade, how vehicles react to obstacles, and whether a patient has a disease—and continuously adapt in response to new data. But these algorithms don't always work smoothly. They don't always make ethical or accurate choices. There are three fundamental reasons for this.

One is simply that the algorithms typically rely on the *probability* that someone will, say, default on a loan or have a disease. Because they make so many predictions, it's likely that *some* will be wrong, just because there's always a chance that they'll be off. The likelihood of

errors depends on a lot of factors, including the amount and quality of the data used to train the algorithms, the specific type of machine-learning method chosen (for example, deep learning, which uses complex mathematical models, versus classification trees that rely on decision rules), and whether the system uses only *explainable algorithms* (meaning humans can describe how they arrived at their decisions), which may not allow it to maximize accuracy.

Second, the environment in which machine learning operates may itself evolve or differ from what the algorithms were developed to face. While this can happen in many ways, two of the most frequent are *concept drift* and *covariate shift*.

With the former the relationship between the inputs the system uses and its outputs isn't stable over time or may be misspecified. Consider a machine-learning algorithm for stock trading. If it has been trained using data only from a period of low market volatility and high economic growth, it may not perform well when the economy enters a recession or experiences turmoil—say, during a crisis like the Covid-19 pandemic. As the market changes, the relationship between the inputs and outputs

—for example, between how leveraged a company is and its stock returns—also may change. Similar misalignment may happen with credit-scoring models at different points in the business cycle.

In medicine, an example of concept drift is when a machine-learning-based diagnostic system that uses skin images as inputs in detecting skin cancers fails to make correct diagnoses because the relationship between, say, the color of someone's skin (which may vary with race or sun exposure) and the diagnosis decision hasn't been adequately captured. Such information often is not even available in electronic health records used to train the machine-learning model.

Idea in Brief

The Problem

Offerings that rely on machine learning are proliferating, raising all sorts of new risks for companies that develop and use them or supply data to train them. That's because such systems don't always make ethical or accurate choices.

The Causes

First, the systems often make decisions based on probabilities. Second, their environments may evolve in an

unanticipated way. Third, their complexity makes it difficult to determine whether or why they made a mistake.

The Solutions

Executives must decide whether to let a system continuously evolve or introduce locked versions at intervals. In addition, they should test the offering appropriately before and after it is rolled out and monitor it constantly once it's on the market.

Covariate shifts occur when the data fed into an algorithm during its use differs from the data that trained it. This can happen even if the patterns the algorithm learned are stable and there's no concept drift. For example, a medical device company may develop its machine-learning-based system using data from large urban hospitals. But once the device is out in the market, the medical data fed into the system by care providers in rural areas may not look like the development data. The urban hospitals might have a higher concentration of patients from certain sociodemographic groups who have underlying medical conditions not commonly seen in rural hospitals. Such disparities may be discovered only when the device makes more errors while out in the market than it did during testing. Given the diversity of markets and the pace at which they're changing, it's

becoming increasingly challenging to foresee what will happen in the environment that systems operate in, and no amount of data can capture all the nuances that occur in the real world.

The third reason machine learning can make inaccurate decisions has to do with the complexity of the overall systems it's embedded in. Consider a device used to diagnose a disease on the basis of images that doctors input—such as IDx-DR, which identifies eye disorders like diabetic retinopathy and macular edema and was the first autonomous machine-learning-based medical device authorized for use by the U.S. Food and Drug Administration. The quality of any diagnosis depends on how clear the images provided are, the specific algorithm used by the device, the data that algorithm was trained with, whether the doctor inputting the images received appropriate instruction, and so on. With so many parameters, it's difficult to assess whether and why such a device may have made a mistake, let alone be certain about its behavior.

But inaccurate decisions are not the only risks with machine learning. Let's look now at two other categories: agency risk and moral risk.

Agency Risk

The imperfections of machine learning raise another important challenge: risks stemming from things that aren't under the control of a specific business or user.

Ordinarily, it's possible to draw on reliable evidence to reconstruct the circumstances that led to an accident. As a result, when one occurs, executives can at least get helpful estimates of the extent of their company's potential liability. But because machine learning is typically embedded within a complex system, it will often be unclear what led to a breakdown—which party, or "agent" (for example, the algorithm developer, the system deployer, or a partner), was responsible for an error and whether there was an issue with the algorithm, with some data fed to it by the user, or with the data used to train it, which may have come from multiple third-party vendors. Environmental change and the probabilistic nature of machine learning make it even harder to attribute responsibility to a particular agent. In fact, accidents or unlawful decisions can occur even without negligence on anyone's part—as there is simply always the possibility of an inaccurate decision.

Executives need to know when their companies are likely to face liability under current law, which may itself also evolve. Consider the medical context. Courts have historically viewed doctors as the final decision-makers and have therefore been hesitant to apply product liability to medical software makers. However, this may change as more black-box or autonomous systems make diagnoses and recommendations without the involvement of (or with much weaker involvement by) physicians in clinics. What will happen, for example, if a machine-learning system recommends a nonstandard treatment for a patient (like a much higher drug dosage than usual) and regulation evolves in such a way that the doctor would most likely be held liable for any harm only if he or she did not follow the system's recommendation? Such regulatory changes may shift liability risks from doctors to the developers of the machine-learning-enabled medical devices, the data providers involved in developing the algorithms, or the companies involved in installing and deploying the algorithms.

Moral Risk

Products and services that make decisions autonomously will also need to resolve ethical dilemmas—a requirement that raises additional risks and regulatory and product development challenges. Scholars have now begun to frame these challenges as problems of *responsible algorithm design*. They include the puzzle of how to automate moral reasoning. Should Tesla, for example, program its cars to think in utilitarian cost-benefit terms or Kantian ones, where certain values cannot be traded off regardless of benefits? Even if the answer is utilitarian, quantification is extremely difficult: How should we program a car to value the lives of three elderly people against, say, the life of one middle-aged person? How should businesses balance trade-offs among, say, privacy, fairness, accuracy, and security? Can all those kinds of risks be avoided?

Moral risks also include biases related to demographic groups. For example, facial-recognition algorithms have a difficult time identifying people of color; skin-lesion-classification systems appear to have unequal accuracy across race; recidivism-prediction instruments give Blacks and Hispanics falsely high ratings, and credit-scoring systems give them unjustly low ones. With many

widespread commercial uses, machine-learning systems may be deemed unfair to a certain group on some dimensions.

The problem is compounded by the multiple and possibly mutually incompatible ways to define fairness and encode it in algorithms. A lending algorithm can be calibrated—meaning that its decisions are independent of group identity after controlling for risk level—while still disproportionately denying loans to creditworthy minorities. As a result, a company can find itself in a "damned if you do, damned if you don't" situation. If it uses algorithms to decide who receives a loan, it may have difficulty avoiding charges that it's discriminating against some groups according to one of the definitions of fairness. Different cultures may also accept different definitions and ethical trade-offs—a problem for products with global markets. A February 2020 European Commission white paper on AI points to these challenges: It calls for the development of AI with "European values," but will such AI be easily exported to regions with different values?

Finally, all these problems can also be caused by model instability. This is a situation where inputs that are close

to one another lead to decisions that are far apart. Unstable algorithms are likely to treat very similar people very differently—and possibly unfairly.

All these considerations, of course, don't mean that we should avoid machine learning altogether. Instead, executives need to embrace the opportunities it creates while making sure they properly address the risks.

To Lock or Not to Lock?

If leaders decide to employ machine learning, a key next question is: Should the company allow it to continuously evolve or instead introduce only tested and locked versions at intervals? Would the latter choice mitigate the risks just described?

This problem is familiar to the medical world. The FDA has so far typically approved only "software as a medical device" (software that can perform its medical functions without hardware) whose algorithms are locked. The reasoning: The agency has not wanted to permit the use of devices whose diagnostic procedures or treatment pathways keep changing in ways it doesn't understand. But as the FDA and other regulators are now realizing,

locking the algorithms may be just as risky, because it doesn't necessarily remove the following dangers:

Inaccurate decisions

Locking doesn't alter the fact that machine-learning algorithms typically base decisions on estimated probabilities. Moreover, while the input of more data usually leads to better performance, it doesn't always, and the amount of improvement can vary; improvements in unlocked algorithms may be greater or smaller for different systems and with different volumes of data. Though it's difficult to understand how the accuracy (or inaccuracy) of decisions may change when an algorithm is unlocked, it's important to try.

Environmental changes

It also matters whether and how the environment in which the system makes decisions is evolving. For example, car autopilots operate in environments that are constantly altered by the behavior of other drivers. Pricing, credit scoring, and trading systems may face a shifting market regime whenever the business cycle enters a new phase. The challenge is ensuring that the

machine-learning system and the environment coevolve in a way that lets the system make appropriate decisions.

Agency risks

Locking an algorithm doesn't eliminate the complexity of the system in which it's embedded. For example, errors caused by using inferior data from third-party vendors to train the algorithm or by differences in skills across users can still occur. Liability can still be challenging to assign across data providers, algorithm developers, deployers, and users.

Moral risks

A locked system may preserve imperfections or biases unknown to its creators. When analyzing mammograms for signs of breast cancer, a locked algorithm would be unable to learn from new subpopulations to which it is applied. Since average breast density can differ by race, this could lead to misdiagnoses if the system screens people from a demographic group that was underrepresented in the training data. Similarly, a credit-scoring algorithm trained on a socioeconomically segregated subset of the population can discriminate

against certain borrowers in much the same way that the illegal practice of redlining does. We want algorithms to correct for such problems as soon as possible by updating themselves as they “observe” more data from subpopulations that may not have been well represented or even identified before. Conversely, devices whose machine-learning systems are not locked could harm one or more groups over time if they’re evolving by using mostly data from a different group. What’s more, identifying the point at which the device gets comparatively worse at treating one group can be hard.

A Tool Kit for Executives

So how should executives manage the existing and emerging risks of machine learning? Developing appropriate processes, increasing the savviness of management and the board, asking the right questions, and adopting the correct mental frame are important steps.

Treat machine learning as if it’s human

Executives need to think of machine learning as a living entity, not an inanimate technology. Just as cognitive testing of employees won't reveal how they'll do when added to a preexisting team in a business, laboratory testing cannot predict the performance of machine-learning systems in the real world. Executives should demand a full analysis of how employees, customers, or other users will apply these systems and react to their decisions. Even when not required to do so by regulators, companies may want to subject their new machine-learning-based products to randomized controlled trials to ensure their safety, efficacy, and fairness prior to rollout. But they may also want to analyze products' decisions in the actual market, where there are various types of users, to see whether the quality of decisions differs across them. In addition, companies should compare the quality of decisions made by the algorithms with those made in the same situations *without* employing them. Before deploying products at scale, especially but not only those that haven't undergone randomized controlled trials, companies should consider testing them in limited markets to get a better idea of their accuracy and behavior when various factors are at

play—for instance, when users don't have equal expertise, the data from sources varies, or the environment changes. Failures in real-world settings signal the need to improve or retire algorithms.

Think like a regulator and certify first

Businesses should develop plans for certifying machine-learning offerings before they go to market. The practices of regulators offer a good road map. In 2019, for example, the FDA published a discussion paper that proposed a new regulatory framework for modifications to machine-learning-based software as a medical device. It laid out an approach that would allow such software to continuously improve while maintaining the safety of patients, which included a complete assessment of the company—or team—developing the software to ensure it had a culture of organizational excellence and high quality that would lead it to regularly test its machine-learning devices. If companies don't adopt such certification processes, they may expose themselves to liability—for example, for performing insufficient due diligence.

Many start-ups provide services to certify that products and processes don't suffer from bias, prejudice,

stereotypes, unfairness, and other pitfalls. Professional organizations, such as the Institute of Electrical and Electronics Engineers and the International Organization for Standardization, are also developing standards for such certification, while companies like Google offer AI ethics services that examine multiple dimensions, ranging from the data used to train systems, to their behavior, to their impact on well-being. Companies might need to develop similar frameworks of their own.

Monitor continuously

As machine-learning-based products and services and the environments they operate in evolve, companies may find that their technologies don't perform as initially intended. It is therefore important that they set up ways to check that these technologies behave within appropriate limits. Other sectors can serve as models. The FDA's Sentinel Initiative draws from disparate data sources, such as electronic health records, to monitor the safety of medical products and can force them to be withdrawn if they don't pass muster. In many ways companies' monitoring programs may be similar to the preventive maintenance tools and processes currently

used by manufacturing or energy companies or in cybersecurity. For example, firms might conduct so-called adversarial attacks on AI like those used to routinely test the strength of IT systems' defenses.

Ask the right questions

Executives and regulators need to delve into the following:

Accuracy and competitiveness. How much is the performance of the machine-learning-based system likely to improve with the volume of new data from its use if we don't lock the algorithm? What will such improvements mean for the business? To what extent will consumers understand the benefits and drawbacks of locked versus unlocked systems?

Biases. What data was used to train the algorithm? How representative is it of the population on which the algorithm will ultimately operate? Can we predict whether an unlocked algorithm will produce less-biased results than a locked one if we allow it to learn over time? Do the algorithm's errors affect minorities or other groups in particular? Can a continuous monitoring approach

establish “guardrails” that stop the algorithm from becoming discriminatory?

The environment. How will the environment in which the offering is used change over time? Are there conditions under which machine learning should not be allowed to make decisions, and if so, what are they? How can we ensure that the offering’s behavior evolves appropriately given how the environment itself is changing? When should we withdraw our offering because the gap between the environment and our offering’s behavior has become too big? What are the boundaries of the environment within which our offering can adapt and operate? How robust and safe are our machine-learning systems throughout their life cycles?

Agency. On which third-party components, including data sources, does the behavior of our machine-learning algorithms depend? How much does it vary when they’re used by different types of people—for example, less-skilled ones? What products or services of other organizations use our data or machine-learning algorithms, possibly exposing us to liability? Should we

allow other organizations to use machine-learning algorithms that we develop?

Develop principles that address your business risks

Businesses will need to establish their own guidelines, including ethical ones, to manage these new risks—as some companies, like Google and Microsoft, have already done. Such guidelines often need to be quite specific (for example, about what definitions of fairness are adopted) to be useful and must be tailored to the risks in question. If you're using machine learning to make hiring decisions, it would be good to have a model that is simple, fair, and transparent. If you're using machine learning to forecast the prices of commodity futures contracts, you may care less about those values and more about the maximum potential financial loss allowed for any decision that machine learning makes.

Luckily, the journey to develop and implement principles doesn't need to be a lonely one. Executives have a lot to learn from the multiyear efforts of institutions such as the OECD, which developed the first intergovernmental AI principles (adopted in 2019 by many countries). The OECD principles promote

innovative, trustworthy, and responsibly transparent AI that respects human rights, the rule of law, diversity, and democratic values, and that drives inclusive growth, sustainable development, and well-being. They also emphasize the robustness, safety, security, and continuous risk management of AI systems throughout their life cycles.

The OECD's recently launched AI Policy Observatory provides further useful resources, such as a comprehensive compilation of AI policies around the world.

Machine learning has tremendous potential. But as this technology, along with other forms of AI, is woven into our economic and social fabric, the risks it poses will increase. For businesses, mitigating them may prove as important as—and possibly more critical than—managing the adoption of machine learning itself. If companies don't establish appropriate practices to address these new risks, they're likely to have trouble gaining traction in the marketplace.

Originally published in January–February 2021. Reprint R2101F

Getting Serious About Diversity

by Robin J. Ely and David A. Thomas

THESE RALLYING CRIES for more diversity in companies, from recent statements by CEOs, are representative of what we hear from business leaders around the world. They have three things in common: All articulate a business case for hiring more women or people of color; all demonstrate good intentions; and none of the claims is actually supported by robust research findings.

We say this as scholars who were among the first to demonstrate the potential benefits of more race and gender heterogeneity in organizations. In 1996 we published an HBR article, "Making Differences Matter: A New Paradigm for Managing Diversity," in which we argued that companies adopting a radically new way of understanding and leveraging diversity could reap the real and full benefits of a diverse workforce. This new way

entailed not only recruiting and retaining more people from underrepresented "identity groups" but also tapping their identity-related knowledge and experiences as resources for learning how the organization could perform its core work better. Our research showed that when companies take this approach, their teams are more effective than either homogeneous teams or diverse teams that don't learn from their members' differences. Such companies send a message that varied points of view are valued and don't need to be suppressed for the sake of group cohesion. This attitude encourages employees to rethink how work gets done and how best to achieve their goals.

We called this approach the *learning-and-effectiveness paradigm*. We argued that cultivating a learning orientation toward diversity—one in which people draw on their experiences as members of particular identity groups to reconceive tasks, products, business processes, and organizational norms—enables companies to increase their effectiveness. We stand by the research on which that article was based, and we continue to advocate its conclusions.

The problem is that nearly 25 years later, organizations have largely failed to adopt a learning orientation toward diversity and are no closer to reaping its benefits. Instead, business leaders and diversity advocates alike are advancing a simplistic and empirically unsubstantiated version of the business case. They misconstrue or ignore what abundant research has now made clear: Increasing the numbers of traditionally underrepresented people in your workforce does not *automatically* produce benefits. Taking an "add diversity and stir" approach, while business continues as usual, will not spur leaps in your firm's effectiveness or financial performance.

And despite all the rhetoric about the value of diversity, white women and people of color remain seriously underrepresented in many industries and in most companies' senior ranks. That lack of progress suggests that top executives don't actually find the business case terribly compelling.

On that point, we have to agree: The *simplistic* business case isn't persuasive. A credible and powerful case *can* be made, however, with three critical modifications. First, platitudes must give way to sound, empirically based conclusions. Second, business leaders must reject the

notion that maximizing shareholder returns is paramount; instead they must embrace a broader vision of success that encompasses learning, innovation, creativity, flexibility, equity, and human dignity. Finally, leaders must acknowledge that increasing demographic diversity does not, by itself, increase effectiveness; what matters is how an organization harnesses diversity, and whether it's willing to reshape its power structure.

In this article we expose the flaws in the current diversity rhetoric and then outline what a 21st-century learning-and-effectiveness paradigm could look like—and how leaders can foster it.

Idea in Brief

The Context

Business leaders often make a business case for diversity, claiming that hiring more women or people of color results in better financial performance.

The Problem

There's no empirical evidence that simply diversifying the workforce, absent fundamental changes to the organizational culture, makes a company more profitable.

A Better Approach

Companies *can* benefit from diversity if leaders create a psychologically safe workplace, combat systems of discrimination and subordination, embrace the styles of employees from different identity groups, and make cultural differences a resource for learning and improving organizational effectiveness.

A Critique of the Business Case for Diversity

Let's start with the claim that putting more women on corporate boards leads to economic gains. That's a fallacy, probably fueled by studies that went viral a decade ago reporting that the more women directors a company has, the better its financial performance. But those studies show correlations, not causality. In all likelihood, some other factor—such as industry or firm size—is responsible for both increases in the number of women directors and improvement in a firm's performance.

In any case, the research touting the link was conducted by consulting firms and financial institutions and fails to pass muster when subjected to scholarly scrutiny. Meta-analyses of rigorous, peer-reviewed studies found no significant relationships—causal or otherwise—between board gender diversity and firm

performance. That could be because women directors may not differ from their male counterparts in the characteristics presumed to affect board decisions, and even if they do differ, their voices may be marginalized. What is more pertinent, however, is that board decisions are typically too far removed from firms' bottom-line performance to exert a direct or unconditional effect.

As for studies citing the positive impact of racial diversity on corporate financial performance, they do not stand up to scrutiny either. Indeed, we know of no evidence to suggest that replacing, say, two or three white male directors with people from underrepresented groups is likely to enhance the profits of a *Fortune* 500 company.

The economic argument for diversity is no more valid when it's applied to changing the makeup of the overall workforce. A 2015 survey of Harvard Business School alumni revealed that 76% of those in senior executive positions believe that "a more diverse workforce improves the organization's financial performance." But scholarly researchers have rarely found that increased diversity leads to improved financial outcomes. They *have* found that it leads to higher-quality work, better decision-making, greater team satisfaction, and more

equality—under certain circumstances. Although those outcomes could conceivably make some aspects of the business more profitable, they would need to be extraordinarily consequential to affect a firm's bottom line.

Moreover, advocates who justify diversity initiatives on the basis of financial benefits may be shooting themselves in the foot. Research suggests that when company diversity statements emphasize the economic payoffs, people from underrepresented groups start questioning whether the organization is a place where they really belong, which reduces their interest in joining it. In addition, when diversity initiatives promise financial gains but fail to deliver, people are likely to withdraw their support for them.

Still another flaw in the familiar business case for diversity is the notion that a diverse team will have richer discussions and a better decision-making process simply because it is diverse. Having people from various identity groups "at the table" is no guarantee that anything will get better; in fact, research shows that things often get worse, because increasing diversity can increase tensions and conflict. Under the right organizational conditions,

though, employees can turn cultural differences into assets for achieving team goals.

Studies have shown, for example, that diverse teams realize performance benefits in certain circumstances: when team members are able to reflect on and discuss team functioning; when status differences among ethnic groups are minimized; when people from both high- and low-status identity groups believe the team supports learning; and—as we reported in our earlier article—when teams orient members to learn from their differences rather than marginalize or deny them. But absent conditions that foster inquiry, egalitarianism, and learning, diversity either is unrelated to or undermines team effectiveness.

Many progressive companies today recognize the conditional nature of the diversity-performance link and have moved beyond "diversity," the catchword of the 1990s, to "diversity and inclusion." They understand that just increasing the number of people from underrepresented groups is not meaningful if those employees do not feel valued and respected. We applaud the emphasis on inclusion, but it is insufficient because it doesn't fundamentally reconfigure power relations.

Being genuinely valued and respected involves more than just feeling included. It involves having the power to help set the agenda, influence what—and how—work is done, have one's needs and interests taken into account, and have one's contributions recognized and rewarded with further opportunities to contribute and advance. Undertaking this shift in power is what the learning-and-effectiveness companies we wrote about in 1996 were doing, and it's what enabled them to tap diversity's true benefits.

The Learning-and-Effectiveness Paradigm, Redux

What we've learned since we wrote our original article is that embracing a learning orientation toward diversity turns out to be quite difficult. To make real progress, people—and the organizational cultures they inhabit—must change. But instead of doing the hard work involved, companies have generally stuck with easier, more limited approaches that don't alter the status quo.

We previously identified four actions that were helping business leaders and managers shift to a learning-and-

effectiveness approach. We still consider those actions fundamental, but we present them anew here to underscore the message in light of today's challenges and opportunities.

Build trust

The first task for those in charge is to build trust by creating a workplace where people feel safe expressing themselves freely. That requires setting a tone of honest discourse and getting comfortable with vulnerability—one's own and others'.

At no time has this need been greater in the United States than during the current unrest spurred by outrage over police brutality against Black men and women—a legacy of centuries of racism. Two weeks into the nationwide protests that began in May, white leaders in companies across the country struggled with how to respond. Publicly expressing support for the Black Lives Matter movement was one thing; knowing what to say to Black employees, who might already have been feeling marginalized or undervalued at work, was quite another. Leaders who were used to wielding authority grounded in their subject-matter expertise had no comparable

expertise to handle the deep grief, rage, and despair felt by many of their employees—especially their Black employees. And Black leaders, many with firsthand experience of police mistreatment and other forms of racial oppression, faced the challenge of managing their own strong emotions and speaking their truth without appearing biased against whites.

Yet troubling times provide opportunities for leaders to begin conversations that foster learning. In response to public acts of racial injustice, for example, white leaders can reach out from a place of vulnerability, as a way of creating connection and psychological safety, rather than staying silent from a place of privilege and self-protection. This was the choice made by a white senior partner in a global professional services firm when he decided to convene a special virtual meeting with his teams across the country. He knew that if he said nothing about the recent racist incidents, his silence would speak for him, with a message not of neutrality but of complicity. Just weeks before, he'd been eloquent in addressing the distress wrought by the Covid-19 pandemic, but when it came to race, he felt at a complete loss. What he astutely realized, though, was that people

needed him simply to begin a dialogue, acknowledge his pain and theirs, and give them the space to talk about their experiences inside and outside the firm, if they wished. He had no solutions, but that moment required none—just a willingness to speak from the heart and listen compassionately to whatever his colleagues might share. Perhaps most important, he was willing to risk not getting his own words or actions exactly right, and he was ready to receive feedback with openness and equanimity.

Actively work against discrimination and subordination

Creating psychological safety and building employees' trust can be an excellent starting point for the second action: taking concrete measures to combat forms of discrimination and subordination that inhibit employees' ability to thrive. This action calls for both individual and collective learning aimed at producing systemic change.

Over the years we've seen the emergence of a multibillion-dollar industry dedicated to advancing such goals. Companies have adopted a slew of initiatives as a result: affinity groups, mentoring programs, work-family accommodation policies, and unconscious-bias training,

to name a few. But the sad truth is that these efforts largely fail to produce meaningful, sustained change—and sometimes even backfire.

Leaders are the stewards of an organization's culture; their behaviors and mindsets reverberate throughout the organization. Hence to dismantle systems of discrimination and subordination, leaders must undergo the same shifts of heart, mind, and behavior that they want for the organization as a whole and then translate those personal shifts into real, lasting change in their companies.

To that end, a first step for leaders is to learn about how systems of privilege and oppression—racism, sexism, ethnocentrism, classism, heterosexism—operate in the wider culture. Numerous excellent books and articles can help with this work; they have the added benefit of relieving those on the receiving end of oppressive systems from the burden of educating their majority-group counterparts. And the impact can be surprising. For example, major news organizations picked up the story of a Black flight attendant who noticed a white male passenger reading a book about white people's reluctance to confront racism. She struck up a conversation with the

man and had a moving exchange with him, eventually learning that he was the CEO of a major airline. The encounter filled her with hope: Here was a powerful executive—someone in a position to effect change—making a genuine effort to understand systemic racism.

Educating oneself is important, but it will be meaningless unless leaders take the next step: investigating how their organization's culture has reproduced systems of oppression, undercutting some groups' opportunities to thrive and succeed, while giving others a boost. As part of that investigation, leaders must examine what stereotypes and assumptions they hold about employees' competencies and suitability for jobs, acknowledge that they have blind spots, and come to see how their personal defenses can shut down learning—their own and their organization's. Working with hundreds of leaders over the years, we have seen how this individual learning journey can be a transformational experience that often leads to individual behavioral change.

But that's not enough. The critical final step in rooting out systems of discrimination and subordination is for leaders to use their personal experience to spur collective

learning and systemic change. It is here that even the most progressive leaders' efforts tend to stall. Such efforts require a well-articulated, widely shared organizational mission to motivate and guide change, together with a collective process of continuous reflection and consciousness-raising, experimentation, and action—followed by sustained attention, monitoring each change for impact, and making adjustments accordingly.

An example of this process comes from a midsize consulting firm whose partners—almost all white men—had begun to fear that high turnover among the white women and people of color they employed meant they were losing talent, potentially undermining the firm's competitiveness. Taking a hard look at their culture, they identified a flawed approach to project assignment that was inadvertently contributing to systematic inequities. Plum projects were going disproportionately to white men; it was the old story of people having an easier time identifying talent when it comes in a package that looks like them. When a particularly challenging project for an important client came up—the kind that can stretch and give exposure to a promising young consultant—the white male partners staffed it with their go-to people:

other white men. Meanwhile, white women and people of color, despite having been recruited from the same highly competitive MBA programs as their white male counterparts, regularly were assigned the more mundane projects. They got stuck doing tasks they had long ago mastered, which led many to leave the firm. Come promotion time, the few who remained were either counseled out or told they still weren't ready for partnership; women waited two years longer than men, on average, to make partner.

But were the go-to people actually better? Did they really have more "raw horsepower," as the partners believed? When those leaders examined their developmental practices, they were chagrined to see clear patterns in who received coaching, whose mistakes were forgiven, and who got second and even third chances to prove themselves: the white men. So after an uncomfortable reckoning with their biases, the partners decided to experiment with making comparable investments in people they'd previously overlooked—people they might have automatically, if not quite consciously, written off simply as hires to meet diversity goals. When they started treating white women and

people of color more like the white men they'd favored, they were surprised to find a bigger, more diverse pool of talent than they'd expected.

Embrace a wide range of styles and voices

The third necessary action for leaders and managers involves actively trying to understand how organizational norms might implicitly discourage certain behavioral styles or silence certain voices. For example, in companies where the prototypical leader is a white man who earns respect by speaking assertively, women and Black men, who are often penalized for being assertive, may find themselves in a double bind: They can conform to the organization's norms and deviate from cultural prescriptions for their group, or they can do the opposite. But either way, they violate one set of expectations, risking marginalization and diminished chances for advancement.

Managers may believe they're giving helpful feedback when they tell a large Black man to smile more so that his white colleagues won't fear him, when they ask a Latina who advocates passionately for a project to dial it down, when they encourage a no-nonsense white woman to be

“nicer,” or when they urge a soft-spoken woman of East Asian descent to speak more forcefully. But all such messages communicate that these employees must be ever-mindful of how others see them in relation to stereotyped images of their group, making it harder for them to bring their talents and perspectives to the table. Companies need performance management systems that tie feedback and evaluation criteria to bona fide task requirements rather than group stereotypes.

Make cultural differences a resource for learning

For companies shifting to a learning-and-effectiveness paradigm, the fourth action is to encourage—and draw lessons from—open discussions about how identity groups shape employees’ experiences inside and outside the organization. Leaders should frame those experiences as a valid source of ideas for enhancing the organization’s work and culture. Even if employees champion ideas that are at odds with the company’s profit goals, those ideas may still be worth pursuing if they help the organization achieve its mission or uphold its values.

Over the years, we have seen that learning from cultural differences is more likely to occur once the

previous three actions are under way: Leaders have created trust, begun to dismantle systems of discrimination and subordination, and embraced a broad range of styles. Without such efforts, talking about differences happens (if it happens at all) only in reaction to diversity-related crises—when discussions tend to be fraught and people's capacity to learn is diminished.

An example of learning from gender diversity comes from Boris Groysberg's study of top-ranked research analysts on Wall Street. In exploring whether they take their star status with them when they switch firms, he found a fascinating sex difference: Unlike their male counterparts, whose performance worsened upon changing firms, women who made a move experienced no such performance drop. The reason, Groysberg concluded, was that women analysts faced sex discrimination, and so they had to do the job differently from men. Women had a more difficult time building support networks inside their firm, had fewer mentors, and were neglected by high-status groups such as the firm's institutional sales force—an important source of industry information. And so, unlike men, women built their franchises on portable, external relationships with

clients, companies, and the media. In addition, they forged unconventional in-house relationships with their firm's retail sales force—also an important source of industry information but a low-status group that male analysts typically ignored. Not only were women stars able to maintain their performance upon switching firms but, generally speaking, they outperformed their male peers over the nine-year period of the study. In short, women were not only different; they were better.

In a follow-up set of case studies, coauthored with Ashish Nanda and Laura Morgan Roberts, respectively, Groysberg showed how a Wall Street firm's research director leveraged women's "difference" to everyone's advantage. He aggressively recruited talented women for the analyst role and then set out to create the conditions that would enable them to thrive, emphasizing team culture, allowing flexible work arrangements, and instituting systems that gave analysts regular, unbiased feedback to help them set personal improvement goals. Additionally, he encouraged people to develop their own style and voice. As one woman star in the firm noted, "We have always been given the freedom to be ourselves." Another said, "I never felt I had to pretend to be male to

fit in here.” Within three years this firm had the highest percentage of top-ranked women analysts of any firm on Wall Street and the lowest rate of female turnover. Furthermore, the research department moved in the rankings from 15th to first, and the unique approach that women had developed for building their franchises became the basis for training all the firm’s analysts. What the research director figured out was that gender had given women analysts a unique set of experiences, and those, together with their resilience and ingenuity, led to new insights into how to do the job better.

We have also seen how the mere act of learning across employees’ differences can have a positive impact, even when the content of the learning is unrelated to people’s identities. The benefits are particularly strong when the differences have been historically fraught with tension. In a study of more than 400 retail bank branches in the northeastern United States, we, together with Irene Padavic of Florida State University, found that the more racially diverse the branch, the better its performance—but only for branches in which *all* employees, across all racial groups, experienced the environment as conducive to learning. Some of that learning definitely came from

sharing cultural knowledge—for example, a white branch manager described how his Chinese coworker's explanations of norms in the Chinese community helped him better serve that segment of customers. But many of the branches' tasks were technical and unrelated to people's cultural backgrounds. In those cases, the benefit from diversity seemed to stem mainly from the process of learning—a process that involves taking risks and being unafraid to say "I don't know," "I made a mistake," or "I need help." Showing such vulnerability across divisive lines of difference, such as race, and being met with acceptance rather than judgment or rejection, strengthens relationships. Stronger relationships in turn increase resilience in the face of conflict and other stressors. In short, for culturally diverse teams, the experience of learning across racial differences can, in and of itself, improve performance.

Inequality is bad for both business and society. Organizations limit their capacity for innovation and continuous improvement unless all employees are full participants in the enterprise: fully seen, heard,

developed, engaged—and rewarded accordingly. Moreover, such treatment can unleash enormous reserves of leadership potential too long suppressed by systems that perpetuate inequality.

When the only legitimate conversation about diversity is one that links it to economic gains, we tend to discount the problem of inequality. In fact, studies have shown that making the economic case diminishes people's sense that equality is itself important, limits socially conscious investors' ability to promote it, and may even increase bias. Furthermore, focusing on financial benefits sends a message to traditionally underrepresented employees that they are worth hiring and investing in only because having "their kind" in the mix increases the firm's profitability.

Companies will not reap benefits from diversity unless they build a culture that insists on equality. Treating differences as a source of knowledge and connection lays the groundwork for such a culture. But as part of that process, firms may have to make financial investments that they won't recoup, at least in the short run, and more will be required of top leaders, managers, and rank-and-file employees alike. Everyone will have to learn how to

actively listen to others' perspectives, have difficult conversations, refrain from blame and judgment, and solicit feedback about how their behaviors and company practices might be impeding the push for a culture that supports learning, equality, and mutual respect. Developing those capacities is no small feat in any context; it is even more challenging for people working across cultural identity differences. But teams that truly embrace the learning-and-effectiveness paradigm will come to understand that homogeneity isn't better; it's just easier. They'll realize, too, that the benefits of diversity arise as much from the collective work of developing those key capacities as from the collective learning they enable.

Finally, while there *is* a business case for diversity—one that rests on sound evidence, an expansive definition of what makes a business successful, and the presence of facilitating conditions—we are disturbed by the implication that there must be economic grounds to justify investing in people from underrepresented groups. Why should anyone need an economic rationale for affirming the agency and dignity of any group of human beings? We should make the necessary investment

because doing so honors our own and others' humanity and gives our lives meaning. If company profits come at the price of our humanity, they are costing us too much. And if diversity initiatives fail to reckon with that trade-off, they will amount to little more than rearranging the deck chairs on a sinking ship.

Originally published in November–December 2020. Reprint R2006J

How to Promote Racial Equity in the Workplace

by Robert Livingston

INTRACTABLE AS IT SEEMS, the problem of racism in the workplace can be effectively addressed with the right information, incentives, and investment. Corporate leaders may not be able to change the world, but they can certainly change *their* world.

Organizations are relatively small, autonomous entities that afford leaders a high level of control over cultural norms and procedural rules, making them ideal places to develop policies and practices that promote racial equity. In this article, I'll offer a practical road map for making profound and sustainable progress toward that goal.

I've devoted much of my academic career to the study of diversity, leadership, and social justice, and over the years I've consulted on these topics with scores of *Fortune* 500

companies, federal agencies, nonprofits, and municipalities. Often, these organizations have called me in because they are in crisis and suffering—they just want a quick fix to stop the pain. But that's akin to asking a physician to write a prescription without first understanding the patient's underlying health condition. Enduring, long-term solutions usually require more than just a pill. Organizations and societies alike must resist the impulse to seek immediate relief for the symptoms, and instead focus on the disease. Otherwise they run the risk of a recurring ailment.

To effectively address racism in your organization, it's important to first build consensus around whether there is a problem (most likely, there is) and, if so, what it is and where it comes from. If many of your employees do not believe that racism against people of color exists in the organization, or if feedback is rising through various communication channels showing that Whites feel that they are the real victims of discrimination, then diversity initiatives will be perceived as the problem, not the solution. This is one of the reasons such initiatives are frequently met with resentment and resistance, often by mid-level managers. Beliefs, not reality, are what

determine how employees respond to efforts taken to increase equity. So, the first step is getting everyone on the same page as to what the reality is and why it is a problem for the organization.

But there's much more to the job than just raising awareness. Effective interventions involve many stages, which I've incorporated into a model I call PRESS. The stages, which organizations must move through sequentially, are: (1) Problem awareness, (2) Root-cause analysis, (3) Empathy, or level of concern about the problem and the people it afflicts, (4) Strategies for addressing the problem, and (5) Sacrifice, or willingness to invest the time, energy, and resources necessary for strategy implementation. Organizations going through these stages move from understanding the underlying condition, to developing genuine concern, to focusing on correction.

A road map for racial equity

Organizations move through these stages sequentially, first establishing an understanding of the underlying condition, then developing genuine concern, and finally focusing on correcting the problem.

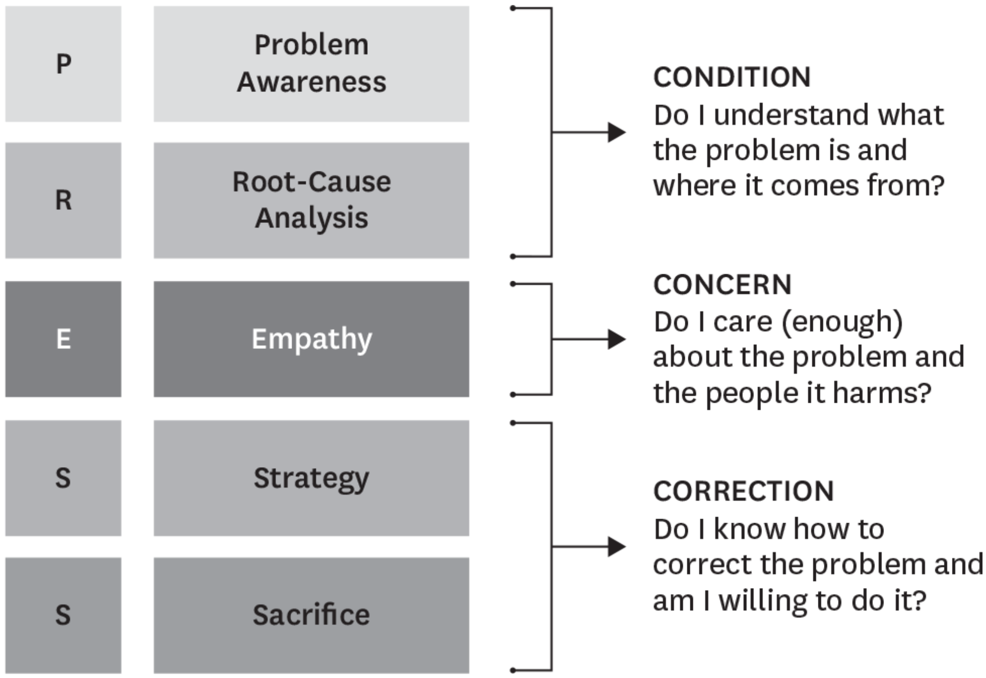

Idea in Brief

The Problem

Racial discrimination—defined as differential evaluation or treatment based solely on race, regardless of intent—remains prevalent in organizations and occurs far more frequently than most White people suspect.

The Opportunity

Intractable as it seems, racism in the workplace can be effectively addressed. Because organizations are autonomous entities that afford leaders a high level of control over norms and policies, they are ideal places to promote racial equity.

The Way Forward

> Effective interventions move through stages, from understanding the underlying condition, to developing genuine concern, to focusing on correction.

Let's now have a closer look at these stages and examine how each informs, at a practical level, the process of working toward racial equity.

Problem Awareness

To a lot of people, it may seem obvious that racism continues to oppress people of color. Yet research consistently reveals that many Whites don't see it that way. For example, a 2011 study by Michael Norton and Sam Sommers found that on the whole, whites in the United States believe that systemic anti-Black racism has steadily decreased over the past 50 years—and that systemic anti-White racism (an implausibility in the United States) has steadily increased over the same time frame. The result: As a group, Whites believe that there is more racism against them than against Blacks. Other recent surveys echo Sommers and Norton's findings, one revealing, for example, that 57% of all Whites and 66% of working-class Whites consider discrimination against whites to be as big a problem as discrimination against Blacks and other people

of color. These beliefs are important, because they can undermine an organization's efforts to address racism by weakening support for diversity policies. (Interestingly, surveys taken since the George Floyd murder indicate an increase in perceptions of systemic racism among Whites. But it's too soon to tell whether those surveys reflect a permanent shift or a temporary uptick in awareness.)

Even managers who recognize racism in society often fail to see it in their own organizations. For example, one senior executive told me, "We don't have any discriminatory policies in our company." However, it is important to recognize that even seemingly "race neutral" policies can enable discrimination. Other executives point to their organizations' commitment to diversity as evidence for the absence of racial discrimination. "Our firm really values diversity and making this a welcoming and inclusive place for everybody to work," another leader remarked.

Despite these beliefs, many studies in the 21st century have documented that racial discrimination is prevalent in the workplace, and that organizations with strong commitments to diversity are no less likely to discriminate. In fact, research by Cheryl Kaiser and colleagues has

demonstrated that the presence of diversity values and structures can actually make matters worse, by lulling an organization into complacency and making Blacks and ethnic minorities more likely to be ignored or harshly treated when they raise valid concerns about racism.

Many White people deny the existence of racism against people of color because they assume that racism is defined by deliberate actions motivated by malice and hatred. However, racism can occur without conscious awareness or intent. When defined simply as differential evaluation or treatment based solely on race, regardless of intent, racism occurs far more frequently than most White people suspect. Let's look at a few examples.

In a well-publicized résumé study by the economists Marianne Bertrand and Sendhil Mullainathan, applicants with White-sounding names (such as Emily Walsh) received, on average, 50% more callbacks for interviews than equally qualified applicants with Black-sounding names (such as Lakisha Washington). The researchers estimated that just being White conferred the same benefit as an additional eight years of work experience—a dramatic head start over equally qualified Black candidates.

Research shows that people of color are well aware of these discriminatory tendencies and sometimes try to counteract them by masking their race. A 2016 study by Sonia Kang and colleagues found that 31% of the Black professionals and 40% of the Asian professionals they interviewed admitted to “Whitening” their résumés, either by adopting a less “ethnic” name or omitting extracurricular experiences (a college club membership, for instance) that might reveal their racial identities.

These findings raise another question: Does Whitening a résumé actually benefit Black and Asian applicants, or does it disadvantage them when applying to organizations seeking to increase diversity? In a follow-up experiment, Kang and her colleagues sent Whitened and non-Whitened résumés of Black or Asian applicants to 1,600 real-world job postings across various industries and geographical areas in the United States. Half of these job postings were from companies that expressed a strong desire to seek diverse candidates. They found that Whitening résumés by altering names and extracurricular experiences increased the callback rate from 10% to nearly 26% for Blacks, and from about 12% to 21% for Asians. What’s particularly unsettling is that a company’s stated commitment to

diversity failed to diminish this preference for Whitened résumés.

This is a very small sample of the many studies that have confirmed the prevalence of racism in the workplace, all of which underscore the fact that people's beliefs and biases must be recognized and addressed as the first step toward progress. Although some leaders acknowledge systemic racism in their organizations and can skip step one, many may need to be convinced that racism persists, despite their "race neutral" policies or pro-diversity statements.

Root-Cause Analysis

Understanding an ailment's roots is critical to choosing the best remedy. Racism can have many psychological sources—cognitive biases, personality characteristics, ideological worldviews, psychological insecurity, perceived threat, or a need for power and ego enhancement. But most racism is the result of structural factors—established laws, institutional practices, and cultural norms. Many of these causes do not involve malicious intent. Nonetheless, managers often misattribute workplace discrimination to the character of individual actors—the so-called bad apples—rather than to broader structural factors. As a result, they

roll out trainings to "fix" employees while dedicating relatively little attention to what may be a toxic organizational culture, for example. It is much easier to pinpoint and blame individuals when problems arise. When police departments face crises related to racism, the knee-jerk response is to fire the officers involved or replace the police chief, rather than examining how the culture licenses, or even encourages, discriminatory behavior.

Appealing to circumstances beyond one's control is another way to exonerate deeply embedded cultural or institutional practices that are responsible for racial disparities. For example, an oceanographic organization I worked with attributed its lack of racial diversity to an insurmountable pipeline problem. "There just aren't any Black people out there studying the migration patterns of the humpback whale," one leader commented. Most leaders were unaware of the National Association of Black Scuba Divers, an organization boasting thousands of members, or of Hampton University, a historically Black college on the Chesapeake Bay, which awards bachelor's degrees in marine and environmental science. Both were entities that could source Black candidates for the job,

especially given that the organization only needed to fill dozens, not thousands, of openings.

A *Fortune* 500 company I worked with cited similar pipeline problems. Closer examination revealed, however, that the real culprit was the culture-based practice of promoting leaders from within the organization—which already had low diversity—rather than conducting a broader industry-wide search when leadership positions became available. The larger lesson here is that an organization's lack of diversity is often tied to inadequate recruitment efforts rather than an empty pipeline. Progress requires a deeper diagnosis of the routine practices that drive the outcomes leaders wish to change.

To help managers and employees understand how being embedded within a biased system can unwittingly influence outcomes and behaviors, I like to ask them to imagine being fish in a stream. In that stream, a current exerts force on everything in the water, moving it downstream. That current is analogous to systemic racism. If you do nothing—just float—the current will carry you along with it, whether you're aware of it or not. If you actively discriminate by swimming with the current, you will be propelled faster. In both cases, the current takes you

in the same direction. From this perspective, racism has less to do with what's in your heart or mind and more to do with how your actions or inactions amplify or enable the systemic dynamics already in place.

Workplace discrimination often comes from well-educated, well-intentioned, open-minded, kindhearted people who are just floating along, severely underestimating the tug of the prevailing current on their actions, positions, and outcomes. Anti-racism requires swimming against that current, like a salmon making its way upstream. It demands much more effort, courage, and determination than simply going with the flow.

In short, organizations must be mindful of the "current," or the structural dynamics that permeate the system, not just the "fish," or individual actors that operate within it.

Empathy

Once people are aware of the problem and its underlying causes, the next question is whether they care enough to do something about it. There is a difference between sympathy and empathy. Many White people experience sympathy, or pity, when they witness racism. But what's more likely to lead to action in confronting the problem is

empathy—experiencing the same hurt and anger that people of color are feeling. People of color want solidarity—and social justice—not sympathy, which simply quiets the symptoms while perpetuating the disease.

One way to increase empathy is through exposure and education. The video of George Floyd's murder exposed people to the ugly reality of racism in a visceral, protracted, and undeniable way. Similarly, in the 1960s, northern Whites witnessed innocent Black protesters being beaten with batons and blasted with fire hoses on television. What best prompts people in an organization to register concern about racism in their midst, I've found, are the moments when their non-White coworkers share vivid, detailed accounts of the negative impact that racism has on their lives. Managers can raise awareness and empathy through psychologically safe listening sessions—for employees who want to share their experiences, without feeling obligated to do so—supplemented by education and experiences that provide historical and scientific evidence of the persistence of racism.

For example, I spoke with Mike Kaufmann, CEO of Cardinal Health—the 16th largest corporation in America—who credited a visit to the Equal Justice Initiative's National Memorial for Peace and Justice, in Montgomery,

Alabama, as a pivotal moment for the company. While diversity and inclusion initiatives have been a priority for Mike and his leadership team for well over a decade, their focus and conversations related to racial inclusion increased significantly during 2019. As he expressed to me, "Some Americans think when slavery ended in the 1860s that African Americans have had an equal opportunity ever since. That's just not true. Institutional systemic racism is still very much alive today; it's never gone away." Kaufmann is planning a comprehensive education program, which will include a trip for executives and other employees to visit the museum, because he is convinced that the experience will change hearts, open eyes, and drive action and behavioral change.

Empathy is critical for making progress toward racial equity because it affects whether individuals or organizations take any action and if so, what kind of action they take. There are at least four ways to respond to racism: join in and add to the injury, ignore it and mind your own business, experience sympathy and bake cookies for the victim, or experience empathic outrage and take measures to promote equal justice. The personal values of individual employees and the core values of the

organization are two factors that affect which actions are undertaken.

Strategy

After the foundation has been laid, it's finally time for the "what do we do about it" stage. Most actionable strategies for change address three distinct but interconnected categories: personal attitudes, informal cultural norms, and formal institutional policies.

To most effectively combat discrimination in the workplace, leaders should consider how they can run interventions on all three of these fronts simultaneously. Focusing only on one is likely to be ineffective and could even backfire. For example, implementing institutional diversity policies without any attempt to create buy-in from employees is likely to produce a backlash. Likewise, focusing just on changing attitudes without also establishing institutional policies that hold people accountable for their decisions and actions may generate little behavioral change among those who don't agree with the policies. Establishing an anti-racist organizational culture, tied to core values and modeled by behavior from the CEO and other top leaders at the company, can

influence both individual attitudes and institutional policies.

Just as there is no shortage of effective strategies for losing weight or promoting environmental sustainability, there are ample strategies for reducing racial bias at the individual, cultural, and institutional levels. The hard part is getting people to actually adopt them. Even the best strategies are worthless without implementation.

I'll discuss how to increase commitment to execution in the final section. But before I do, I want to give a specific example of an institutional strategy that works. It comes from Massport, a public organization that owns Boston Logan International Airport and commercial lots worth billions of dollars. When its leaders decided they wanted to increase diversity and inclusion in real estate development in Boston's booming Seaport District, they decided to leverage their land to do it. Massport's leaders made formal changes to the selection criteria determining who is awarded lucrative contracts to build and operate hotels and other large commercial buildings on their parcels. In addition to evaluating three traditional criteria—the developer's experience and financial capital, Massport's revenue potential, and the project's architectural design—

they added a fourth criterion called "comprehensive diversity and inclusion," which accounted for 25% of the proposal's overall score, the same as the other three. This forced developers not only to think more deeply about how to create diversity but also to go out and do it. Similarly, organizations can integrate diversity and inclusion into managers' scorecards for raises and promotions—if they think it's important enough. I've found that the real barrier to diversity is not figuring out "What can we do?" but rather "Are we willing to do it?"

Sacrifice

Many organizations that desire greater diversity, equity, and inclusion may not be willing to invest the time, energy, resources, and commitment necessary to make it happen. Actions are often inhibited by the assumption that achieving one desired goal requires sacrificing another desired goal. But that's not always the case. Although nothing worth having is completely free, racial equity often costs less than people may assume. Seemingly conflicting goals or competing commitments are often relatively easy to reconcile—once the underlying assumptions have been identified.

As a society, are we sacrificing public safety and social order when police routinely treat people of color with compassion and respect? No. In fact, it's possible that kinder policing will actually increase public safety. Famously, the city of Camden, New Jersey, witnessed a 40% drop in violent crime after it reformed its police department, in 2012, and put a much greater emphasis on community policing.

The assumptions of sacrifice have enormous implications for the hiring and promotion of diverse talent, for at least two reasons. First, people often assume that increasing diversity means sacrificing principles of fairness and merit, because it requires giving "special" favors to people of color rather than treating everyone the same.

People often assume that fairness means treating everyone *equally*, or exactly the same. In reality, fairness requires treating people equitably—which may entail treating people differently, but in a way that makes sense.

Of course, what is "sensible" depends on the context and the perceiver. Does it make sense for someone with a physical disability to have a parking space closer to a building? Is it fair for new parents to have six weeks of paid leave to be able to care for their baby? Is it right to allow

active-duty military personnel to board an airplane early to express gratitude for their service? My answer is yes to all three questions, but not everyone will agree. For this reason, equity presents a greater challenge to gaining consensus than equality.

In thinking about fairness in the context of American society, leaders must consider the unlevel playing fields and other barriers that exist—provided they are aware of systemic racism. They must also have the courage to make difficult or controversial calls. For example, it might make sense to have an employee resource group for Black employees but not white employees. Fair outcomes may require a process of treating people differently. To be clear, different treatment is not the same as "special" treatment—the latter is tied to favoritism, not equity.

One leader who understands the difference is Maria Klawe, the president of Harvey Mudd College. She concluded that the only way to increase the representation of women in computer science was to treat men and women differently. Men and women tended to have different levels of computing experience prior to entering college—different levels of *experience*, not intelligence or potential. Society treats boys and girls differently

throughout secondary school—encouraging STEM subjects for boys but liberal arts subjects for girls, creating gaps in experience. To compensate for this gap created by bias in society, the college designed two introductory computer-science tracks—one for students with no computing experience and one for students with some computing experience in high school. The no-experience course tended to be 50% women whereas the some-experience course was predominantly men. By the end of the semester, the students in both courses were on par with one another. Through this and other equity-based interventions, Klawe and her team were able to dramatically increase the representation of women and minority computer-science majors and graduates.

The second assumption many people have is that increasing diversity requires sacrificing high quality and standards. Consider again the fence scenario. All three people have the same height or "potential." What varies is the level of the field and the fence—apt metaphors for privilege and discrimination, respectively. Because the person on the far left has lower barriers to access, does it make sense to treat the other two people differently to compensate? Do we have an obligation to do so when

differences in outcomes are caused by the field and the fence, not someone's height? Maria Klawe sure thought so. How much human potential is left unrealized within organizations because we do not recognize the barriers that exist?

Finally, it's important to understand that quality is difficult to measure with precision. There is no test, instrument, survey, or interviewing technique that will enable you to invariably predict who the "best candidate" will be. The NFL draft illustrates the difficulty in predicting future job performance: Despite large scouting departments, plentiful video of prior performance, and extensive tryouts, almost half of first round picks turn out to be busts. This may be true for organizations as well. Research by Sheldon Zedeck and colleagues on corporate hiring processes has found that even the best screening or aptitude tests predict only 25% of intended outcomes, and that candidate quality is better reflected by "statistical bands" rather than a strict rank ordering. This means that there may be absolutely no difference in quality between the candidate who scored first out of 50 people and the candidate who scored eighth.

The big takeaway here is that "sacrifice" may actually involve giving up very little. If we look at people within a band of potential and choose the diverse candidate (for example, number eight) over the top scorer, we haven't sacrificed quality at all—statistically speaking—even if people's intuitions lead them to conclude otherwise.

Managers should abandon the notion that a "best candidate" must be found. That kind of search amounts to chasing unicorns. Instead, they should focus on hiring well-qualified people who show good promise, and then should invest time, effort, and resources into helping them reach their potential.

The tragedies and protests we have witnessed this year across the United States have increased public awareness and concern about racism as a persistent problem in our society. The question we now must confront is whether, as a nation, we are willing to do the hard work necessary to change widespread attitudes, assumptions, policies, and practices. Unlike society at large, the workplace very often requires contact and cooperation among people from different racial, ethnic, and cultural backgrounds. Therefore, leaders should host open and candid

conversations about how their organizations are doing at each of the five stages of the model—and use their power to *press* for profound and perennial progress.

Originally published in September–October 2020. Reprint R2005D

Our Work-from-Anywhere Future

by Prithwiraj (Raj) Choudhury

BEFORE 2020 A MOVEMENT was brewing within knowledge-work organizations. Personal technology and digital connectivity had advanced so far and so fast that people had begun to ask, "Do we really need to be together, in an office, to do our work?" We got our answer during the pandemic lockdowns. We learned that a great many of us don't in fact need to be colocated with colleagues on-site to do our jobs. Individuals, teams, entire workforces, can perform well while being entirely distributed—and they have. So now we face new questions: Are all-remote or majority-remote organizations the future of knowledge work? Is work from anywhere (WFA) here to stay?

Without question, the model offers notable benefits to companies and their employees. Organizations can reduce or eliminate real estate costs, hire and use talent

globally while mitigating immigration issues, and, research indicates, perhaps enjoy productivity gains. Workers get geographic flexibility (that is, live where they prefer to), eliminate commutes, and report better work/life balance. However, concerns persist regarding how WFA affects communication, including brainstorming and problem-solving; knowledge sharing; socialization, camaraderie, and mentoring; performance evaluation and compensation; and data security and regulation.

To better understand how leaders can capture the upside of WFA while overcoming the challenges and avoiding negative outcomes, I've studied several companies that have embraced all- or majority-remote models. They include the United States Patent and Trademark Office, or USPTO (which has several thousand WFA workers); Tulsa Remote; Tata Consultancy Services, or TCS (a global IT services company that has announced a plan to be 75% remote by 2025); GitLab (the world's largest all-remote company, with 1,300 employees); Zapier (a workflow automation company with more than 300 employees, none of them colocated, around the

United States and in 23 other countries); and MobSquad (a Canadian start-up that employs WFA workers).

The Covid-19 crisis has opened senior leaders' minds to the idea of adopting WFA for all or part of their workforces. In addition to TCS, companies including Twitter, Facebook, Shopify, Siemens, and State Bank of India have announced that they will make remote work permanent even after a vaccine is available. Another organization I've studied is BRAC, one of the world's largest NGOs, which is headquartered in Bangladesh. Forced into remote work this year, it is deciding what work model to adopt for the long term.

If your organization is considering a WFA program, transition, or launch, this article can provide a guide.

A Short History of Remote Work

A large-scale transition from traditional, colocated work to remote work arguably began with the adoption of work-from-home (WFH) policies in the 1970s, as soaring gasoline prices caused by the 1973 OPEC oil embargo made commuting more expensive. Those policies allowed people to eschew physical offices in favor of their homes,

coworking spaces, or other community locations, such as coffee shops and public libraries, for occasional days, on a regular part-time basis, or full-time, with the expectation that they would come into the office periodically. Workers were often also given control over their schedules, allowing them to make time for school pickups, errands, or midday exercise without being seen as shirking. They saved time by commuting less and tended to take fewer sick days.

Thanks to the advent of personal computers, the internet, email, broadband connectivity, laptops, cell phones, cloud computing, and videotelephony, the adoption of WFH increased in the 2000s. As the researchers Ravi S. Gajendran and David A. Harrison note in a 2007 article, this trend was accelerated by the need to comply with, for example, the Americans with Disabilities Act of 1990 and mandates of the U.S. Equal Employment Opportunity Commission.

Idea in Brief

The Shift

The Covid-19 lockdowns proved that it is not only possible but perhaps preferable for knowledge workers to do their jobs

from anywhere. Will this mark a long-term shift into all-remote work?

Benefits and Challenges

Studies show that working from home yields numerous benefits for both individuals and their organizations, most notably in the form of enhanced productivity and engagement. But when all or most employees are remote, challenges arise for communication, knowledge sharing, socialization, performance evaluation, security, and more.

The Research

As more companies adopt work-from-anywhere policies, best practices are emerging. The experiences of GitLab, Tata Consultancy Services, Zapier, and others show how the risks associated with this type of work can be overcome.

Research has shown performance benefits. A 2015 study by Nicholas Bloom and coauthors found that when employees opted in to WFH policies, their productivity increased by 13%. When, nine months later, the same workers were given a choice between remaining at home and returning to the office, those who chose the former saw even further improvements: They were 22% more productive than they had been before the experiment. This suggests that people should probably determine for

themselves the situation (home or office) that fits them best.

In recent years many companies have allowed more employees to work from home. It's true that several prominent corporations, including Yahoo and IBM, had reversed course before the pandemic, asking their employees to resume colocated work in a bid to spur more-effective collaboration. But other organizations—the ones I study—moved toward greater geographic flexibility, allowing some if not all employees, new and old, to work from anywhere, completely untethered to an office. The USPTO is a prime example. Its leaders launched a WFA program in 2012, building on an existing WFH program that mandated workers' physical presence at headquarters, in northern Virginia, at least one day a week. The WFA program, in contrast, requires employees to spend two years at HQ followed by a WFH phase, after which they may live anywhere in the continental United States, provided they're willing to pay out of pocket for periodic travel back to headquarters (totaling no more than 12 days a year). The patent examiners in the program dispersed all across the country, choosing to move closer

to family, to better climates, or to places with a lower cost of living.

Most companies that offer WFH or WFA options keep some workers—at the USPTO it's trainees and administrators—at one or more offices. In other words, they are hybrid-remote operations. But the experiment with all-remote work forced by Covid-19 has caused some of these organizations to strategically move toward majority-remote, with fewer than 50% of employees colocated in physical offices. TCS, for example, which employs close to 418,000 people who were traditionally located either on campuses or at client sites around the world, has decided to adopt a 25/25 model: Employees will spend only 25% of their working hours in the office, and at no point will the company have more than 25% of workers colocated. TCS plans to complete this transition in five years.

Even before the crisis, a smaller group of companies had taken this trend a step further, eliminating offices altogether and dispersing everyone, from entry-level associates to the CEO. GitLab embraces this model at scale: Its remote workers span sales, engineering,

marketing, personnel management, and executive roles in more than 60 countries.

Exploring the Benefits

I've spent the past five years studying the practices and productivity trends of WFA companies. The upsides—for individuals, companies, and society—are clear. Let me outline them.

For individuals

One striking finding is how greatly workers benefit from these arrangements. Many told me that they regard the freedom to live anywhere in the world as an important plus. For those in dual-career situations, it eases the pain of looking for two jobs in a single location. One patent examiner told me, "I'm a military spouse, which means I live in a world with frequent moves and personal upheavals that prevent many spouses from pursuing lasting careers of their choice. WFA has been the most meaningful telework program I have encountered. It allows me to follow my husband to any U.S. state at a

moment’s notice and pursue my own aspirations to contribute to my home and society.”

Some cited a better quality of life. “WFA has allowed my children to see their grandparents on a regular basis and play with their cousins,” I heard from another USPTO employee. “Being closer to family has improved my overall happiness.” Others talked about proximity to medical care for children, accommodating their partners, and the ability to enjoy warmer weather, prettier views, and greater recreational opportunities. Millennials in particular seemed captivated by the idea that WFA would allow them to become “digital nomads,” traveling the world while still employed. Before the pandemic-related restrictions, some companies, such as Remote Year, were aiming to facilitate that lifestyle, and some countries, such as Estonia and Barbados, have created a new class of employment visa for such workers. As one patent examiner said, “Participation in [WFA] is outstanding for work/life balance. I live in my favorite part of the country . . . I have more time to relax.”

Cost of living was another frequent theme. Because the USPTO did not adjust salaries according to where employees chose to live, one patent examiner told me, “I

was able to buy a large home in my new location for about a quarter of the cost in northern Virginia.” Some localities, such as the state of Vermont and the city of Tulsa, Oklahoma (where Tulsa Remote is located), have made a concerted effort to lure remote workers, touting the local community and lower costs. In San Francisco the average rent on a two-bedroom apartment is $4,128; in Tulsa it’s a mere $675.

WFA also helps knowledge workers deal with immigration issues and other restrictions on their ability to secure good jobs. William Kerr, Susie Ma, and I recently studied MobSquad, whose coworking spaces in Halifax, Calgary, and other cities enable talented knowledge workers to bypass the onerous U.S. visa and green card system and instead obtain fast-track work permits from Canada’s Global Talent Stream. Thus they can continue serving companies and clients in the United States and other countries while living and paying taxes in Canada.

One engineer we interviewed had come to the United States after graduating from high school in his home country at the age of 12. At age 16 he enrolled at a U.S. university, where he acquired degrees in math, physics, and computer science in three years. By age 19 he was

employed at a medical tech company through the optional practical training (OPT) program, but he failed to get an H-1B visa and faced deportation. MobSquad moved him to Calgary, and he kept working with the same employer.

In interviews with female employees at BRAC, I learned that women whose careers were previously limited by cultural taboos against traveling to remote places or delegating housework had been helped by WFA. As one explained, "Earlier I had to wake up early in the morning and cook three meals for my intergenerational family. Working remotely has allowed me to spread out the household work, get extra sleep, and be more productive."

For organizations

My research also uncovered ample organizational benefits from WFA programs. For example, they increase employee engagement—an important metric of success for any company. In 2013, a year after it instituted work from anywhere, the USPTO was ranked highest on the Best Places to Work in the Federal Government survey.

Workers are not only happier but also more productive. When Cirrus Foroughi, Barbara Larson, and I evaluated the USPTO's transition from WFH to WFA, the timing of which happened at random for workers who'd chosen that path, we found that WFA boosted individual productivity by 4.4%, as measured by the number of patents examined each month. The switch also led examiners to exert greater effort. Of course, further research is needed to determine whether WFA generates similar benefits for workers performing different tasks in other team structures and organizations.

Some gains generated by WFA are more obvious. For example, fewer in-office employees means smaller space requirements and reduced real estate costs. The USPTO estimated that increases in remote work in 2015 saved it $38.2 million. WFA programs also hugely expand an organization's potential talent pool to include workers tied to a location far from that of the company. That's a primary reason for the adoption at TCS of what it calls secure borderless workspaces, or SBWs: It wants to ensure that every project is staffed by employees with the right skills, no matter where they are. Rajesh Gopinathan, the CEO, describes this model as "talent on the cloud," while

another senior executive says it will potentially allow the company to tap niche labor markets, such as Eastern Europe, that have a large supply of skilled financial analysts and data scientists.

Finally, WFA can reduce attrition. Some USPTO workers explained that because they loved their preferred locales but also recognized the limited job opportunities there, they were motivated to work harder and stay longer with the Patent Office. Leaders at GitLab, too, pointed to employee retention as a positive outcome of the company's decision to be all-remote. The net benefit, they believe, including the productivity increases and property cost savings they've seen, equals $18,000 a year for each worker.

For society

WFA organizations have the potential to reverse the brain drain that often plagues emerging markets, small towns, and rural locations. In fact, Tulsa Remote was established to attract diverse, energetic, community-minded newcomers to a city still healing from historic race riots a century ago. With an offer of $10,000 to relocate to Tulsa, the company attracted more than 10,000 applications for

just 250 slots from 2019 to 2020. Obum Ukabam was one of the workers chosen. When he's not busy with his day job as a marketing manager, he mentors and coaches a local high school debate team. Talented newcomers of varied ethnicities are arguably making the city more multicultural. Meanwhile, the transitions at the USPTO and TCS have brought many people back to their hometowns.

Remote work helps the environment as well. In 2018 Americans' commute time averaged 27.1 minutes each way, or about 4.5 hours a week. Eliminating that commute—particularly in places where most people commute by car—generates a significant reduction in emissions. The USPTO estimates that in 2015 its remote workers drove 84 million fewer miles than if they had been traveling to headquarters, reducing carbon emissions by more than 44,000 tons.

Addressing the Concerns

The office—with its meeting rooms and break areas and opportunities for both formal and informal interaction—has been a way of life for so long that it's hard to imagine

getting rid of it. And legitimate hurdles exist to making all-remote work not only manageable but successful. However, the Covid-19 all-remote experiment has taught many knowledge-work organizations and their employees that with time and attention, those concerns can be addressed. And in the companies I've studied, several best practices are emerging.

Communication, brainstorming, and problem-solving

When workers are distributed, synchronous communication becomes more difficult. Tools such as Zoom, Skype, Microsoft Teams, and Google Hangouts can help for those working in the same or similar time zones but not for those spread farther apart. In research with Jasmina Chauvin and Tommy Pan Fang, I found that when changing to or from daylight saving time caused a one- to two-hour reduction in business-hour overlap (BHO) between offices of a very large global corporation, the volume of communication fell by 9.2%, primarily among production workers. When BHO was greater, R&D staffers conducted more unplanned synchronous calls. Group meetings are even harder to schedule. Nadia

Vatalidis of GitLab's People Operations group says that having team members in Manila, Nairobi, Johannesburg, Raleigh, and Boulder made finding a time for their weekly group call nearly impossible.

WFA organizations must therefore get comfortable with asynchronous communication, whether through a Slack channel, a customized intracompany portal, or even a shared Google document in which geographically distributed team members write their questions and comments and trust that other team members in distant time zones will respond at the first opportunity. One benefit to this approach is that employees are more likely to share early-stage ideas, plans, and documents and to welcome early feedback; the pressure to present polished work is less than it would be in more formal, synchronous team meetings. GitLab calls this process blameless problem-solving. The company's leaders note that employees accustomed to a culture of emails, phone calls, and meetings may struggle to change old habits; they solve that problem with training during onboarding and beyond. At Zapier, in a program called Zap Pal, each new hire is matched with an experienced buddy who sets up at least one introductory Zoom call and continues to

check in throughout the first month. For synchronous brainstorming the company uses video calls and online whiteboards such as Miro, Stormboard, IPEVO Annotator, Limnu, and MURAL but also urges employees to use asynchronous means of problem-solving through Slack channel threads.

Knowledge sharing

This is another challenge for all-remote or majority-remote organizations. Distributed colleagues can't tap one another on the shoulder to ask questions or get help. Research by Robin Cowan, Paul David, and Dominique Foray has postulated that much workplace knowledge is not codified (even when it can be) and instead resides "in people's heads." This is a problem for all organizations, but much more so for those that have embraced WFA. The companies I've studied solve it with transparent and easily accessible documentation. At GitLab all team members have access to a "working handbook," which some describe as "the central repository for how we run the company." It currently consists of 5,000 searchable pages. All employees are encouraged to add to it and taught how to create a new topic page, edit an existing

one, embed video, and so forth. Ahead of meetings, organizers post agendas that link to the relevant sections to allow invitees to read background information and post questions. Afterward recordings of the sessions are posted on GitLab's YouTube channel, agendas are edited, and the handbook is updated to reflect any decisions made.

Employees may see the extra work of documentation as a "tax" and balk at the extremely high level of transparency necessary for a WFA organization to thrive. Thorsten Grohsjean and I have argued that senior managers must set an example on these fronts by codifying knowledge and freely sharing information while explaining that these are necessary trade-offs to allow for geographic flexibility.

A related idea is to create transcripts, publicly post slides, and record video seminars, presentations, and meetings to create a repository of such material that individuals can view asynchronously at their convenience. For its 2020 annual meeting, which was forced by the pandemic to go virtual, the Academy of Management curated 1,120 prerecorded sessions, arguably expanding the flow of knowledge to scholars—

especially those in emerging markets—far more than would have been possible at the in-person event, which typically happens in North America.

Socialization, camaraderie, and mentoring

Another major worry, cited by managers and workers alike, is the potential for people to feel isolated socially and professionally, disconnected from colleagues and the company itself, particularly in organizations where some people are colocated and some are not. Research by Cecily D. Cooper and Nancy B. Kurland has shown that remote workers often feel cut off from the information flow they would typically get in a physical office. Without in-person check-ins, managers may miss signs of growing burnout or team dysfunction. Even with videoconferencing that allows for reading body language and facial expressions, the concern is that virtual colleagues are less likely to become close friends because their face-to-face interactions are less frequent. As GitLab's technical evangelist Priyanka Sharma put it, "I was very nervous when I was first thinking of joining, because I was very social in the office. I worried that I would be so lonely at home and wouldn't have that

community feel." Houda Elyazgi, a marketing executive on the Tulsa Remote team, expressed similar sentiments: "Remote work can be very isolating, especially for introverts. You almost have to create an intentional experience when you're socializing with others. And then you have to be 'on' all the time, even when you're trying to relax. That's taxing."

In my research I've seen a range of policies that seek to address these concerns and create opportunities for socialization and the spreading of company norms. Many WFA organizations rely on technology to help facilitate virtual watercoolers and "planned randomized interactions," whereby someone in the company schedules groups of employees to chat online. Some use AI and virtual reality tools to pair up remote colleagues for weekly chats. For example, Sike Insights is using data on individual communication styles and AI to create Slackbot buddies, while eXp Realty, an all-remote company I'm currently researching, uses a VR platform called VirBELA to create a place for distant team members to gather in avatar form.

Sid Sijbrandij, a cofounder and the CEO of GitLab, told me, "I know at Pixar they placed the restroom centrally

so people would bump into each other—but why depend on randomness for that? Why not step it up a notch and actually organize the informal communication?” These “mixers” often include senior and C-suite executives. When I described them to my HBS colleague Christina Wallace, she gave them a nice name: *community collisions*. And companies have always needed to manufacture them: Research dating back to Thomas J. Allen’s work at MIT in the 1970s shows that workers colocated on the same “campus” may not experience serendipitous interactions if they are separated by a wall, a ceiling, or a building.

When it comes to interaction between people at different hierarchical levels, my research has revealed two problems with straightforward solutions. Iavor Bojinov, Ashesh Rambachan, and I found that the senior leaders of a global firm were often too stretched to offer one-on-one mentoring to virtual workers. So we implemented a Q&A process whereby workers posed questions through a survey and leaders responded asynchronously. Senior managers at another global firm told me that they had difficulty being themselves on camera. Whereas young remote workers were “living

their lives on Instagram," their older colleagues found virtual engagement harder. The company implemented coaching sessions to make those executives more comfortable on Microsoft Teams.

Another solution to the socialization problem is to host "temporary colocation events," inviting all workers to spend a few days with colleagues in person. Prior to Covid-19, Zapier hosted two of those a year, paying for employee flights, accommodation, and food; giving teams an activities budget; and sending people home with $50 to use on a thank-you gift for their loved ones. Carly Moulton, the company's senior communications specialist, told me, "Personally, I have made a lot of friendships with the people I travel to and from the airport with. The event managers will put us into random groups based on what time you arrive and depart. I've always been with people I don't normally work with, so it's nice to have a dedicated time when you have to make conversation."

Finally, at the USPTO, I learned another way to create camaraderie. Several WFA examiners have voluntarily created "remote communities of practice" so that a handful of them can get together periodically. A group

living in North Carolina, for example, decided to schedule meetings on a golf course to socialize, discuss work, and problem-solve together. Another manager created a “virtual meal” by ordering the same pizza for delivery to the homes of all remote direct reports during a weekly team call.

Performance evaluation and compensation

How can you rate and review employees you’re never physically with, particularly on “soft” but important metrics such as interpersonal skills? All-remote companies evaluate remote workers according to the quality of their work output, the quality of virtual interactions, and feedback from clients and colleagues. Zapier, for example, uses Help Scout for customer support replies; a feature of this software is that customers can submit a “happiness score” by rating the response as “great,” “OK,” or “not good.”

In the spring and summer of 2020, as groups suddenly transitioned to remote work, I was asked whether managers should use software to track worker productivity and prevent shirking. I am very much opposed to this Orwellian approach. The USPTO

addressed claims of “examiner fraud” and “attendance abuse” in its WFA program following a review by the U.S. Commerce Department’s Office of the Inspector General. Those claims involved either overreporting of hours worked or shifts in the time logs of completed work, such as backloading at the end of a calendar quarter—neither of which related to the metric on which performance was measured: the number of patents examined. Nevertheless, from then on, all USPTO teleworkers had to use organizational IT tools, such as logging in to a virtual private network (VPN), having a presence indicator turned on, and using the same messaging services. But when we compared data from before and after that intervention, we found that it had no effect on average output.

How to set compensation for workers who work from anywhere is an active and interesting debate. As mentioned, it’s a benefit to be able to reside in a lower-cost-of-living locale while earning the income one would in a more expensive one. But that’s conditional on the company’s not adjusting wages according to where a worker lives, as was the case at the USPTO. Matt Mullenweg, the founder of Automattic (parent of

WordPress), another all-remote company, told me that its policy is to pay the same wages for the same roles, regardless of location. But GitLab and other companies do have different pay for different geographies, taking into account the experience of the worker, the contract type, and the task being performed. Although research is needed on which approach is optimal, it's possible that companies that tie wages to location will lose high-quality WFA workers to rivals that don't. Another pertinent issue is whether to pay WFA workers in the currency of the country where the organization is incorporated or the local one, in part to ensure consistent wages across locations over time given exchange-rate fluctuations.

Data security and regulation

Several managers told me that cybersecurity was a big area of focus for WFA programs and organizations. "What if the WFA worker takes photographs of client data screens and sends them to a competitor?" one asked. The CIOs of some companies with remote-work policies said another key concern was employees' use of personal, less-protected devices for work at home.

It's true that all-remote companies have to work harder to protect employee, corporate, and customer data. As TCS transitions to a majority-remote model, it has moved from "perimeter-based security" (whereby the IT team attempts to secure every device) to "transaction-based security" (whereby machine learning algorithms analyze any abnormal activities on any employee laptop). MobSquad has replicated its client security infrastructure for WFA workers, and employees work on clients' cloud, email, and hardware in its offices for security reasons. All-remote and majority-remote organizations I have studied are experimenting with a wide range of solutions to protect client data using predictive analytics, data visualization, and computer vision.

Transitioning to an all-remote or a majority-remote organization sometimes requires jumping regulatory hurdles as well. At the onset of the pandemic, when TCS was forced to become all-remote, it had to work with NASSCOM (India's National Association of Software and Service Companies) and the Indian authorities to change laws overnight so that call center staffers could work from home. Other laws had to be tweaked so that TCS workers could take laptops and other equipment out of physical

offices located in India’s “special economic zones.” Irfhan Rawji, the founder and CEO of MobSquad, had to work closely with the Canadian government to ensure that the economic migrants chosen by the company to move to Canada could receive their expedited work permits and be integrated into its model. Any all-remote organization thinking about hiring talent globally has to consider local labor laws as they relate to hiring, compensation, pensions, vacation, and sick leave.

Is This Right for Your Organization?

Of course, WFA may not be possible at this time for some organizations, such as manufacturing companies—though that could change with advances in 3D printing, automation, digital twins, and other technologies. However, with the right strategy, organizational processes, technologies, and—most important—leadership, many more companies, teams, and functions than one might have thought could go all or mostly remote. My ongoing research with Jan Bena and David Rowat suggests, for example, that start-up knowledge-work companies, particularly in the tech sector, are well

positioned to adopt a WFA model from their inception. Take the all-remote eXp Realty: We found that lower real estate, utility, and other overhead costs may mean a higher valuation for the company if and when its founders exit the start-up.

My studies of the USPTO and TCS indicate that large and mature organizations, too, can successfully transition to a hybrid or a majority-remote regime. The question is not whether work from anywhere is possible but what is needed to make it possible. The short answer: management. "If all senior leaders are working from an office, then workers would be drawn to that location to get face time," one all-remote middle manager told me. But if leaders support synchronous and asynchronous communication, brainstorming, and problem-solving; lead initiatives to codify knowledge online; encourage virtual socialization, team building, and mentoring; invest in and enforce data security; work with government stakeholders to ensure regulatory compliance; and set an example by becoming WFA employees themselves, all-remote organizations may indeed emerge as the future of work.

Originally published in November–December 2020. Reprint
R2006C

A More Sustainable Supply Chain

by Verónica H. Villena and Dennis A. Gioia

IN RECENT YEARS a rising number of multinational corporations have pledged to work only with suppliers that adhere to social and environmental standards. Typically, these MNCs expect their first-tier suppliers to comply with those standards, and they ask that those suppliers in turn ask for compliance from *their* suppliers—who ideally ask the same from *their* suppliers. And so on. The aim is to create a cascade of sustainable practices that flows smoothly throughout the supply chain, or, as we prefer to call it, the supply network.

It's an admirable idea, but it's been hard to realize in practice. Many of the MNCs that have committed to it have faced scandals brought about by suppliers that, despite being aware of sustainability standards, have nevertheless gone on to violate them. Consider the embarrassing

scrutiny that Apple, Dell, and HP endured not long ago for sourcing electronics from overseas companies that required employees to work in hazardous conditions, and the fallout that Nike and Adidas suffered for using suppliers that were dumping toxins into rivers in China.

What's more, all those scandals involved first-tier suppliers. The practices of lower-tier suppliers are almost always worse, increasing companies' exposure to serious financial, social, and environmental risks. In this article we describe various ways that MNCs can defuse the ticking time bomb those risks represent.

Where the Problems Are

To understand the situation and develop ideas for tackling it, we conducted a study of three supply networks. Each was headed by an MNC considered to be a "sustainability leader"—one in the automotive industry, one in electronics, and one in pharmaceuticals and consumer products. (For the specific selection criteria, see the "About the Research" sidebar.) We also studied a representative set of each MNC's suppliers—a total of nine top-tier and 22 lower-tier suppliers, based variously in Mexico, China, Taiwan, and the United States. What we discovered was

that many were violating the standards that the MNCs expected them to adhere to. The hoped-for cascade effect was seldom occurring.

About the Research

WE FOCUSED OUR STUDY on three "exemplary" multinational corporations that met five selection criteria: (1) They were included in the Dow Jones Sustainability Index. (2) They were members of the Carbon Disclosure Project (CDP) and the United Nations Global Compact. (3) They had been involved in industrywide supply-chain sustainability efforts. (4) They were certified as having a large percentage of plants with effective quality-management systems (ISO 9001), environmental management systems (ISO 14001), and safety-management systems (OHSAS 18001). (5) They were members of the Billion Dollar Roundtable (firms spending at least $1 billion with minority- and women-owned suppliers).

We also interviewed representatives of industry associations (including the Responsible Business Alliance and the Automotive Industry Action Group) and NGOs (including the CDP and the Centre for Reflection and Action on Labour Rights) to gain a more comprehensive view of how each of these stakeholders helps MNCs disseminate their sustainability agendas throughout their supply networks.

For more information about the research, see "The Missing Link? The Strategic Role of Procurement in Building Sustainable Supply Networks," by Verónica H. Villena, *Production and Operations Management* (May 2019), and "On the Riskiness of

Lower-Tier Suppliers: Managing Sustainability in Supply Networks," by Verónica H. Villena and Dennis A. Gioia, *Journal of Operations Management* (November 2018).

We found problems in every country we studied. In Mexico we visited five lower-tier suppliers; all lacked environmental management systems, and four lacked procedures for handling red-flag social problems such as sexual harassment, retaliation by supervisors, and hazardous labor conditions. At three of the companies, temporary workers made up nearly 50% of the workforce, and turnover rates sometimes reached 100%, making it difficult to implement viable environmental, health, and safety programs. In China and Taiwan we visited 10 lower-tier suppliers, all of which had marginal environmental practices, dangerous working conditions, and chronic overtime issues. In the United States we studied seven lower-tier suppliers and found that three had high concentrations of airborne chemicals and a lack of systematic accident reporting.

Idea in Brief

The Problem

Many multinational corporations have committed themselves to using suppliers with sustainable social and environmental

practices, but suppliers—especially those low in the supply chain—often don't comply with standards. This poses serious financial, social, and environmental risks.

The Research

The authors studied the supply networks of three MNCs considered to be sustainability leaders. They discovered a set of best practices—but also saw how difficult it can be to enforce standards.

The Solution

Awareness is key. Companies should consider adopting the best practices featured in this article, such as establishing long-term sustainability goals and including lower-tier suppliers in an overall sustainability strategy.

The pattern is worrisome. Remember, all those suppliers were connected to model firms that were working proactively to encourage sustainability. If exemplary MNCs are having trouble ensuring good practices among their lower-tier suppliers, then "regular" firms, in all likelihood, are faring even worse at this.

The problem, ironically, often starts with the MNCs themselves. They frequently place orders that exceed suppliers' capacity or impose unrealistic deadlines, leading supplier factories to demand heavy overtime from their workers. When we asked a representative at one supplier

why his company had violated a 60-hour workweek limit, he gave us a frank explanation: “We didn’t want to tell our customer that we can’t produce its products on time, because otherwise it’s going to try to find someone else that can. But our customer didn’t give us enough notice to hire enough skilled people to do the job.”

First-tier suppliers, for their part, rarely concern themselves with their own suppliers’ sustainability practices. That’s often because they’re struggling with sustainability issues themselves. The noncompliant company we cited above, for example, doesn’t try to enforce a strict 60-hour workweek limit with any of its suppliers. “We don’t comply with this requirement ourselves,” the representative told us, “so how could we ask our own suppliers to do so?”

For MNCs, there are special challenges in governing lower-tier suppliers. There’s often no direct contractual relationship, and a particular MNC’s business often doesn’t mean that much to the lower-tier supplier. If American and Japanese automakers rely heavily on a certain seat maker, for example, they can demand that it adhere to their sustainability standards. But that seat maker may have a hard time getting *its* suppliers to follow suit. Suppose it

does business with a foam manufacturer that has many other big customers in the electronics, appliance, and health care industries—each of which has different sustainability standards. The foam manufacturer has little incentive to conform to the automakers' sustainability requirements, because the automakers account for only a small fraction of its total business.

Furthermore, most lower-tier suppliers are not well known, so they receive relatively little attention and pressure from the media, NGOs, and other stakeholders. Even when they do attract attention (for sexual harassment problems, for example, or chronic overtime demands), we found that they do not feel the need to address the issues involved. They tend to act only when MNCs intervene.

Lower-tier suppliers are also the least equipped to handle sustainability requirements. They often do not have sustainability expertise or resources, and they may be unaware of accepted social and environmental practices and regulations. They are also frequently located in countries where such regulations are nonexistent, lax, or not enforced at all. And typically they don't know much about the sustainability requirements imposed by MNCs—but even if they do, they have no incentive to comply. This

may explain why most of the lower-tier suppliers in our study lacked programs to dispose of toxic waste and in fact had no environmental management program whatsoever.

MNCs, too, are handicapped by ignorance. They frequently don't even know who their lower-tier suppliers are, let alone where they're located or what capabilities they have (or don't have). Many of the 22 lower-tier suppliers in our study are small or medium-size private firms that provide little information to the public—characteristics that, in effect, make them almost invisible. Several directors of the three MNCs we studied viewed this as a big problem. "The demon in this place," one of them said, "is the [lower-tier] suppliers that I know the least about." Another said, "I don't have control over the ones that pose the highest risks, so I'm losing sleep over them."

All these concerns mean that lower-tier suppliers are unquestionably the riskiest members of a supply network. If they have poor or dubious sustainability performance, then an MNC that does business with them can endanger its reputation and suffer profound repercussions—losing customers, being forced to find new suppliers, or having its supply chain disrupted. To reduce such risks, MNCs need to

include both first-tier and lower-tier suppliers in their sustainability programs.

Best Practices

The three MNCs in our study have taken a number of steps to promote suppliers' social and environmental responsibility:

- They have established long-term sustainability goals.
- They require first-tier suppliers to set their own long-term sustainability goals.
- They include lower-tier suppliers in the overall sustainability strategy.
- They task a point person on staff with extending the firm's sustainability program to first- and lower-tier suppliers.

These are all beneficial measures that other companies should consider adopting. Firms can also borrow some of the specific strategies that our MNCs use to spread good practices throughout their supply networks. (See the exhibit "Managing lower-tier supplier sustainability.") These fall into four broad categories:

Managing lower-tier supplier sustainability

Ideally, multinational corporations will use a combination of approaches—direct, indirect, collective, and global—to encourage sustainable practices throughout their supply networks. Some specific strategies within each type of approach are listed below.

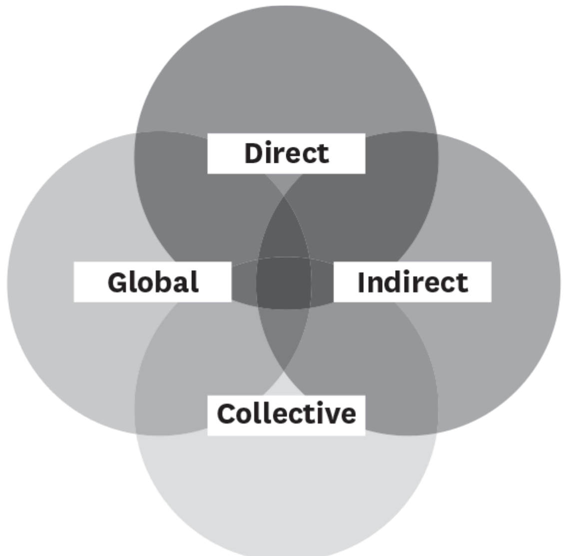

Direct

- Evaluate first-tier suppliers by using sustainability performance indicators that capture their requirements for lower-tier suppliers.
- Survey suppliers on their environmental, health, safety, and labor practices and on their procurement practices.
- Work with major first-tier suppliers to map the firm's supply network.

Indirect

- Provide training and foster peer learning among first-tier suppliers to help them improve their procurement practices with lower-tier suppliers.
- Select high-performing suppliers to pilot new sustainability initiatives.
- Reward suppliers for cascading sustainability requirements to lower-tier suppliers.

Collective

- Commit to developing and complying with industrywide sustainability standards, and help suppliers become full members of industry organizations.
- Via industry organizations, share resources with competitors and major suppliers to achieve sustainability goals.
- Encourage first- and lower-tier suppliers to take advantage of sustainability training programs offered by industry organizations.

Global

- Work closely with relevant NGOs and international institutions interested in improving supply chain sustainability.
- Use tools and data that those organizations provide for dealing with suppliers (contracts and scorecards).
- Recognize suppliers that excel in programs sponsored by NGOs and international institutions.

Direct approach

The MNCs we studied set and monitor social and environmental targets for their first-tier suppliers regarding second-tier suppliers. The automotive corporation, for instance, has a strong commitment to supplier diversity. It requires its first-tier suppliers to allocate 7% of their procurement spending to minority suppliers. Some first-tier suppliers were already meeting that target; others have made substantial changes to do so (for example, by changing performance criteria for their purchasing managers). The first-tier suppliers we interviewed noted that the MNC periodically checks to see if the target is being met and creates opportunities to help them network with minority lower-tier suppliers.

Another MNC annually surveys its first-tier suppliers to gather information not only about their health, safety, labor, and environmental practices but also about the sustainability performance of their lower-tier suppliers. The surveys seem to be having the desired effect: They've prompted first-tier suppliers to engage in internal discussions about whether they should and could alter their procurement practices (to adopt industrywide

sustainability standards, for example). And on two occasions, firms have made changes to comply with MNC requirements (such as using key performance indicators to monitor supplier sustainability).

Additionally, the three MNCs work with their major suppliers to map the connections and interdependencies in their supply networks, including those at the lower-tier level. This allows them to identify potentially risky lower-tier suppliers and to work with the major suppliers to deploy customized risk-mitigation programs where needed.

Indirect approach

The MNCs we studied delegate elements of lower-tier-supplier sustainability management to their first-tier suppliers. This approach is effective because the MNCs are hands-on: They offer training to suppliers and provide some incentives for implementing sustainability practices. Most of the first-tier suppliers we interviewed told us that such training had led them to make substantial changes in their manufacturing processes and to begin asking *their* suppliers to adopt similar sustainability standards.

The three MNCs have also created preferred-supplier programs aimed at fostering peer learning about

sustainability. One corporation, for instance, invites its most socially and environmentally responsible suppliers to join an exclusive group that enables them to strengthen relationships with the MNC and exchange best sustainability practices with one another. Several of these suppliers have started to set their own sustainability requirements for the suppliers they use.

To further encourage first-tier suppliers to cascade the MNCs' sustainability requirements into their own supply networks, MNCs can use supplier sustainability awards, long-term contracts, and preferred status.

Collective approach

Our MNCs collaborate with their competitors and major suppliers to develop and disseminate industrywide sustainability standards. They recognize that a single MNC cannot be expected to fight alone against the problematic labor or environmental practices of global suppliers. Doing so would be not only prohibitively expensive but also unfair, because in most sectors, the major corporations use many of the same suppliers.

The MNCs we studied are all founding members of industry associations focused on developing sustainability standards, providing assessment tools, and offering

training to first- and lower-tier suppliers. One notable association is the Responsible Business Alliance (RBA), whose members include Intel, HP, IBM, Dell, Philips, and Apple.

Collaborative initiatives have many benefits. They can increase efficiencies for suppliers, who can use a standardized self-assessment or audit to satisfy many customers and thus avoid duplication. These initiatives can also draw in more suppliers, because suppliers that have many customers with the same sustainability requirements tend to be more willing to participate. And collaboration can make sustainability initiatives more feasible, because industrywide training is subsidized by members.

Additionally, when MNCs help their first-tier suppliers become full members of an industry association, those suppliers must then comply with industry standards, which means they have to assess their own suppliers' sustainability. The RBA, for example, requires its full members to conduct approved audits annually for at least 25% of their own high-risk facilities *and* 25% of their high-risk suppliers' facilities. (Risk here is assessed along labor, health and safety, environmental, and ethical dimensions.)

Industry associations have a unique power over both first- and lower-tier suppliers, as most of their members are major players in their sectors. Consider the electronics maker Flex, a full member of the RBA and a first-tier supplier for many MNCs. A second-tier electronics supplier is unlikely to refuse a request from Flex for a compliance audit, because it knows that Flex itself has passed this audit and that most other top-tier electronics suppliers, to stay competitive, will probably start issuing similar audit requests.

Global approach

The MNCs we studied make a point of collaborating with international organizations and NGOs that share their goals. For instance, all three corporations have joined the United Nations Global Compact, an international effort to promote corporate social responsibility. The three MNCs also participate in the Carbon Disclosure Project's (CDP's) Supply Chain Program, a global data-collection platform in which suppliers disclose information about their carbon emissions. Firms such as Microsoft, Johnson & Johnson, and Walmart use this platform to engage their suppliers in being transparent about their environmental impact. Several participating suppliers told us that as a result, they

are now collecting previously unsolicited information and making investments to try to reduce their carbon footprints.

The progress is encouraging: According to the CDP's 2019 supply chain report, 35% of the program members engaged with their suppliers on climate change in 2018, up from 23% the year before. Additionally, the report noted, "as suppliers become more mature in their understanding of sustainability issues and advance their approaches for taking action, there is evidence that they too are improving in their efforts to cascade positive change downwards through their own supply chains." This is occurring not only because MNCs have asked their suppliers to disclose their carbon emissions but also because that information influences how the MNCs contract with suppliers. One of the corporations we studied has created an award to recognize the suppliers that have improved the most in terms of CDP Supply Chain Program performance. Another MNC includes the program's ratings in its supplier scorecard and monitors those ratings annually.

Room for Improvement

The MNCs in our study have successfully addressed some of the problematic sustainability practices of their suppliers. But as we've already noted, there's plenty of room for improvement in what they're doing. In our research, we identified a few critical shortcomings in their operations when it comes to developing sustainability beyond first-tier suppliers.

First, the MNCs' engineering and procurement units often preapprove lower-tier suppliers, but their vetting criteria don't include social and environmental considerations. In other words, engineering and procurement address only the first of the proverbial three Ps of sustainability (profit), focusing on such issues as cost, quality, delivery, and technology, while overlooking the second and third Ps (people and the planet). Not surprisingly, that can lead to situations in which preapproved lower-tier suppliers violate the sustainability requirements of the MNCs they work with. The first-tier suppliers are then in a tough spot. Like it or not, they have to work with preapproved suppliers—but they are held accountable if those companies mistreat workers or harm the environment. As one exasperated manager said while

describing this conundrum, “I am just using the supplier you asked me to use!”

Such predicaments are not uncommon. Different functional units of an MNC (engineering, procurement, sustainability) may pursue different agendas in interacting with first- and lower-tier suppliers—with results that do systemic damage to the corporation’s overall sustainability effort and undermine its credibility. To avoid this, MNCs should set convergent sustainability goals and align the incentives for *all* functions that interact with first- and lower-tier suppliers.

A second problem is lack of sustainability training and incentives for procurement officers. All of the 52 procurement employees we interviewed (at MNCs and at suppliers) said they needed more training to properly pursue supplier sustainability on behalf of their firms. Arguably, they need more incentives as well: Companies must reward them for hitting all three Ps—that is, not just cost, quality, and delivery goals but also social and environmental ones. Our research suggests that isn’t yet happening in a meaningful way. For the procurement professionals we interviewed, cost savings were unquestionably the top priority, followed by quality

improvement and on-time delivery. Social and environmental concerns were notably absent. We should add that although companies at every level of the supply network need to provide more training and incentives for their procurement officers, supplier firms are likely to do so only if MNCs lead the way.

A third shortcoming we observed is that although our three MNCs devote considerable effort to developing their first-tier suppliers' sustainability capabilities, they have little direct contact with their first-tier suppliers' procurement personnel. As a result, those people are poorly informed about the MNCs' sustainability requirements and cannot communicate them clearly to their own suppliers, much less enforce them. To alleviate that problem, MNCs could invite suppliers' procurement personnel to their sustainability training sessions (along with environmental, health, and safety personnel) and encourage them to participate in industrywide sustainability training. Alternatively, MNCs could engage the top executives at their first-tier suppliers and explain the importance of building a sustainable supply network, with the goal of motivating them to catalyze the

dissemination of sustainability requirements to lower-tier suppliers.

Many multinational corporations sincerely want to embed fair labor practices and environmental responsibility throughout their supply networks. A good way to start is by adopting the sustainability strategies used by the three MNCs in our study. But all corporations can and should do more. They should send their suppliers a more consistent message that economic, social, and environmental requirements are *all* important. They should make the same message clear to their procurement officials and create incentives for them to pursue not only economic goals but also environmental and social goals. Those officials should take a hands-on approach to collecting data about suppliers' capacity, monitoring indicators of their sustainability performance, and engaging with them in continuous improvement projects. The MNCs should also work directly with their suppliers' procurement units on the best ways to disseminate sustainability requirements throughout their supply networks. The danger of not acting is clear: A supply chain is only as strong as its weakest link.

Originally published in March–April 2020. Reprint R2002F

How Apple Is Organized for Innovation

by Joel M. Podolny and Morten T. Hansen

APPLE IS WELL KNOWN FOR ITS innovations in hardware, software, and services. Thanks to them, it grew from some 8,000 employees and $7 billion in revenue in 1997, the year Steve Jobs returned, to 137,000 employees and $260 billion in revenue in 2019. Much less well known are the organizational design and the associated leadership model that have played a crucial role in the company's innovation success.

When Jobs arrived back at Apple, it had a conventional structure for a company of its size and scope. It was divided into business units, each with its own P&L responsibilities. General managers ran the Macintosh products group, the information appliances division, and the server products division, among others. As is often the case with

decentralized business units, managers were inclined to fight with one another, over transfer prices in particular. Believing that conventional management had stifled innovation, Jobs, in his first year returning as CEO, laid off the general managers of all the business units (in a single day), put the entire company under one P&L, and combined the disparate functional departments of the business units into one functional organization. (See the exhibit "Apple's functional organization.")

Apple's functional organization

In 1997, when Steve Jobs returned to Apple, it had a conventional structure for its size and scope. It was divided into business units, each with its own P&L responsibilities. After retaking the helm, Jobs put the entire company under one P&L and combined the disparate departments of the business units into one functional organization that aligns expertise with decision rights—a structure Apple retains to this day.

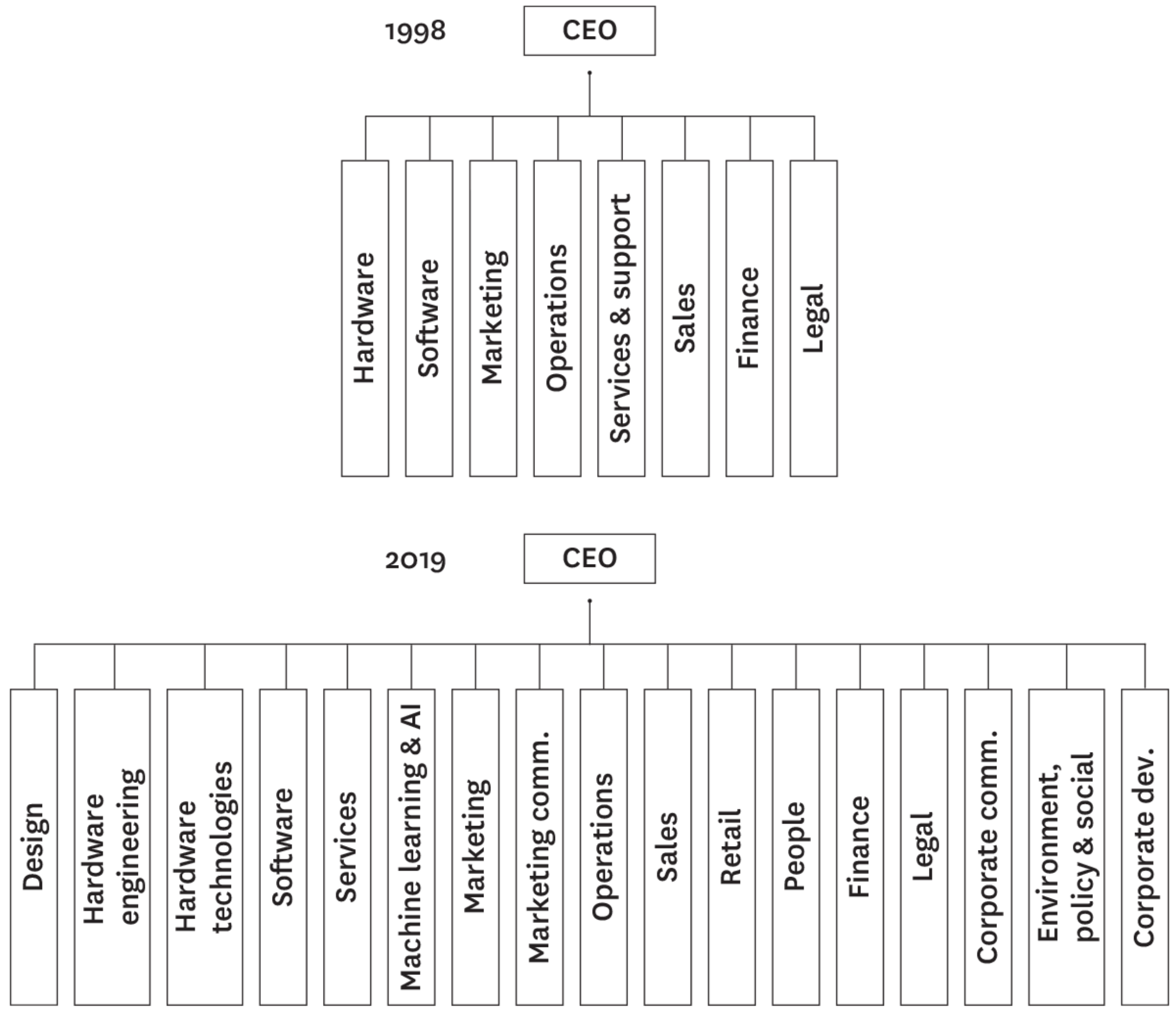

The adoption of a functional structure may have been unsurprising for a company of Apple's size at the time. What *is* surprising—in fact, remarkable—is that Apple retains it today, even though the company is nearly 40 times as large in terms of revenue and far more complex than it was in 1998. Senior vice presidents are in charge of functions, not products. As was the case with Jobs before him, CEO Tim Cook occupies the only position on the

organizational chart where the design, engineering, operations, marketing, and retail of any of Apple's main products meet. In effect, besides the CEO, the company operates with no conventional general managers: people who control an entire process from product development through sales and are judged according to a P&L statement.

Idea in Brief

The Challenge

Major companies competing in many industries struggle to stay abreast of rapidly changing technologies.

One Major Cause

They are typically organized into business units, each with its own set of functions. Thus the key decision makers—the unit leaders—lack a deep understanding of all the domains that answer to them.

The Apple Model

The company is organized around functions, and expertise aligns with decision rights. Leaders are cross-functionally collaborative and deeply knowledgeable about details.

Business history and organizational theory make the case that as entrepreneurial firms grow large and complex, they must shift from a functional to a multidivisional structure

to align accountability and control and prevent the congestion that occurs when countless decisions flow up the org chart to the very top. Giving business unit leaders full control over key functions allows them to do what is best to meet the needs of their individual units' customers and maximize their results, and it enables the executives overseeing them to assess their performance. As the Harvard Business School historian Alfred Chandler documented, U.S. companies such as DuPont and General Motors moved from a functional to a multidivisional structure in the early 20th century. By the latter half of the century the vast majority of large corporations had followed suit. Apple proves that this conventional approach is not necessary and that the functional structure may benefit companies facing tremendous technological change and industry upheaval.

Apple's commitment to a functional organization does not mean that its structure has remained static. As the importance of artificial intelligence and other new areas has increased, that structure has changed. Here we discuss the innovation benefits and leadership challenges of Apple's distinctive and ever-evolving organizational model, which may be useful for individuals and companies

wanting to better understand how to succeed in rapidly changing environments.

Why a Functional Organization?

Apple's main purpose is to create products that enrich people's daily lives. That involves not only developing entirely new product categories such as the iPhone and the Apple Watch, but also continually innovating within those categories. Perhaps no product feature better reflects Apple's commitment to continuous innovation than the iPhone camera. When the iPhone was introduced, in 2007, Steve Jobs devoted only six seconds to its camera in the annual keynote event for unveiling new products. Since then iPhone camera technology has contributed to the photography industry with a stream of innovations: High dynamic range imaging (2010), panorama photos (2012), True Tone flash (2013), optical image stabilization (2015), the dual-lens camera (2016), portrait mode (2016), portrait lighting (2017), and night mode (2019) are but a few of the improvements.

To create such innovations, Apple relies on a structure that centers on functional expertise. Its fundamental belief is that those with the most expertise and experience in a

domain should have decision rights for that domain. This is based on two views: First, Apple competes in markets where the rates of technological change and disruption are high, so it must rely on the judgment and intuition of people with deep knowledge of the technologies responsible for disruption. Long before it can get market feedback and solid market forecasts, the company must make bets about which technologies and designs are likely to succeed in smartphones, computers, and so on. Relying on technical experts rather than general managers increases the odds that those bets will pay off.

Second, Apple's commitment to offer the best possible products would be undercut if short-term profit and cost targets were the overriding criteria for judging investments and leaders. Significantly, the bonuses of senior R&D executives are based on companywide performance numbers rather than the costs of or revenue from particular products. Thus product decisions are somewhat insulated from short-term financial pressures. The finance team is not involved in the product road map meetings of engineering teams, and engineering teams are not involved in pricing decisions.

We don't mean to suggest that Apple doesn't consider costs and revenue goals when deciding which technologies and features the company will pursue. It does, but in ways that differ from those employed by conventionally organized companies. Instead of using overall cost and price targets as fixed parameters within which to make design and engineering choices, R&D leaders are expected to weigh the benefits to users of those choices against cost considerations.

In a functional organization, individual and team reputations act as a control mechanism in placing bets. A case in point is the decision to introduce the dual-lens camera with portrait mode in the iPhone 7 Plus in 2016. It was a big wager that the camera's impact on users would be sufficiently great to justify its significant cost.

One executive told us that Paul Hubel, a senior leader who played a central role in the portrait mode effort, was "out over his skis," meaning that he and his team were taking a big risk: If users were unwilling to pay a premium for a phone with a more costly and better camera, the team would most likely have less credibility the next time it proposed an expensive upgrade or feature. The camera turned out to be a defining feature for the iPhone 7 Plus,

and its success further enhanced the reputations of Hubel and his team.

It's easier to get the balance right between an attention to costs and the value added to the user experience when the leaders making decisions are those with deep expertise in their areas rather than general managers being held accountable primarily for meeting numerical targets. Whereas the fundamental principle of a conventional business unit structure is to align accountability and control, the fundamental principle of a functional organization is to align expertise and decision rights.

Thus the link between how Apple is organized and the type of innovations it produces is clear. As Chandler famously argued, "structure follows strategy"—even though Apple doesn't use the structure that he anticipated large multinationals would adopt.

Now let's turn to the leadership model underlying Apple's structure.

Three Leadership Characteristics

Ever since Steve Jobs implemented the functional organization, Apple's managers at every level, from senior vice president on down, have been expected to possess

three key leadership characteristics: deep expertise that allows them to meaningfully engage in all the work being done within their individual functions; immersion in the details of those functions; and a willingness to collaboratively debate other functions during collective decision-making. When managers have these attributes, decisions are made in a coordinated fashion by the people most qualified to make them.

Deep expertise

Apple is not a company where general managers oversee managers; rather, it is a company where experts lead experts. The assumption is that it's easier to train an expert to manage well than to train a manager to be an expert. At Apple, hardware experts manage hardware, software experts software, and so on. (Deviations from this principle are rare.) This approach cascades down all levels of the organization through areas of ever-increasing specialization. Apple's leaders believe that world-class talent wants to work for and with other world-class talent in a specialty. It's like joining a sports team where you get to learn from and play with the best.

Early on, Steve Jobs came to embrace the idea that managers at Apple should be experts in their area of

management. In a 1984 interview he said, “We went through that stage in Apple where we went out and thought, *Oh, we’re gonna be a big company, let’s hire professional management.* We went out and hired a bunch of professional management. It didn’t work at all. . . . They knew how to manage, but they didn’t know how to *do* anything. If you’re a great person, why do you want to work for somebody you can’t learn anything from? And you know what’s interesting? You know who the best managers are? They are the great individual contributors who never, ever want to be a manager but decide they have to be . . . because no one else is going to . . . do as good a job.”

One current example is Roger Rosner, who heads Apple’s software application business, which includes work-productivity apps such as Pages (word processing), Numbers (spreadsheets), and Keynote (presentations) along with GarageBand (music composition), iMovie (movie editing), and News (an app providing news content). Rosner, who studied electrical engineering at Carnegie Mellon, joined Apple in 2001 as a senior engineering manager and rose to become the director of iWork applications, the vice president of productivity apps, and since 2013 the VP of applications. With his deep

expertise gained from previous experience as the director of engineering at several smaller software companies, Rosner exemplifies an expert leading experts.

In a functional organization, experts leading experts means that specialists create a deep bench in a given area, where they can learn from one another. For example, Apple's more than 600 experts on camera hardware technology work in a group led by Graham Townsend, a camera expert. Because iPhones, iPads, laptops, and desktop computers all include cameras, these experts would be scattered across product lines if Apple were organized in business units. That would dilute their collective expertise, reducing their power to solve problems and generate and refine innovations.

Immersion in the details

One principle that permeates Apple is "Leaders should know the details of their organization three levels down," because that is essential for speedy and effective cross-functional decision-making at the highest levels. If managers attend a decision-making meeting without the details at their disposal, the decision must either be made without the details or postponed. Managers tell war stories about making presentations to senior leaders who drill

down into cells on a spreadsheet, lines of code, or a test result on a product.

Of course, the leaders of many companies insist that they and their teams are steeped in the details. But few organizations match Apple. Consider how its senior leaders pay extreme attention to the exact shape of products' rounded corners. The standard method for rounding corners is to use an arc of a circle to connect the perpendicular sides of a rectangular object, which produces a somewhat abrupt transition from straight to curve. In contrast, Apple's leaders insist on continuous curves, resulting in a shape known in the design community as a "squircle": The slope starts sooner but is less abrupt. (See the exhibit "One example of Apple's attention to detail.") An advantage of hardware products without abrupt changes in curvature is that they produce softer highlights (that is, little to no jump in light reflection along the corner). The difference is subtle, and executing on it isn't simply a matter of a more complicated mathematical formula. It demands that Apple's operations leaders commit to extremely precise manufacturing tolerances to produce millions of iPhones and other products with squircles. This deep immersion in detail isn't just a concern

that is pushed down to lower-level people; it is central at the leadership level.

One example of Apple's attention to detail

The standard method for rounding the corners of a rectangular object is to use an arc of a circle to connect the object's perpendicular sides. That can result in an abrupt transition in curvature. To produce softer highlights by minimizing light reflection, Apple uses a "squircle," which creates continuous curves.

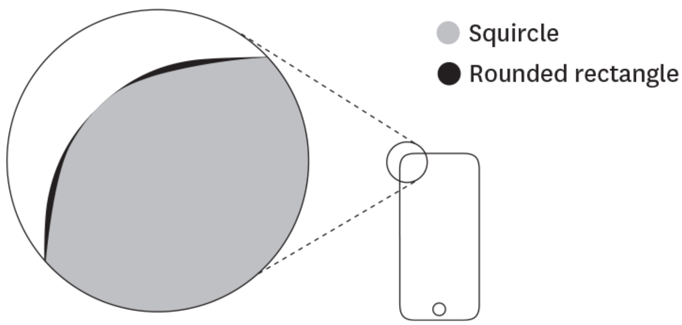

Source: Apple

Having leaders who are experts in their areas and can go deep into the details has profound implications for how Apple is run. Leaders can push, probe, and "smell" an issue. They know which details are important and where to focus their attention. Many people at Apple see it as liberating, even exhilarating, to work for experts, who provide better guidance and mentoring than a general

manager would. Together, all can strive to do the best work of their lives in their chosen area.

Willingness to collaboratively debate

Apple has hundreds of specialist teams across the company, dozens of which may be needed for even one key component of a new product offering. For example, the dual-lens camera with portrait mode required the collaboration of no fewer than 40 specialist teams: silicon design, camera software, reliability engineering, motion sensor hardware, video engineering, core motion, and camera sensor design, to name just a few. How on earth does Apple develop and ship products that require such coordination? The answer is collaborative debate. Because no function is responsible for a product or a service on its own, cross-functional collaboration is crucial.

When debates reach an impasse, as some inevitably do, higher-level managers weigh in as tiebreakers, including at times the CEO and the senior VPs. To do this at speed with sufficient attention to detail is challenging for even the best of leaders, making it all the more important that the company fill many senior positions from within the ranks of its VPs, who have experience in Apple's way of operating.

However, given Apple's size and scope, even the executive team can resolve only a limited number of stalemates. The many horizontal dependencies mean that ineffective peer relationships at the VP and director levels have the potential to undermine not only particular projects but the entire company. Consequently, for people to attain and remain in a leadership position within a function, they must be highly effective collaborators.

That doesn't mean people can't express their points of view. Leaders are expected to hold strong, well-grounded views and advocate forcefully for them, yet also be willing to change their minds when presented with evidence that others' views are better. Doing so is not always easy, of course. A leader's ability to be both partisan and open-minded is facilitated by two things: deep understanding of and devotion to the company's values and common purpose, and a commitment to separating how *right* from how *hard* a particular path is so that the difficulty of executing a decision doesn't prevent its being selected.

The development of the iPhone's portrait mode illustrates a fanatical attention to detail at the leadership level, intense collaborative debate among teams, and the power of a shared purpose to shape and ultimately resolve

debates. In 2009 Hubel had the idea of developing an iPhone feature that would allow people to take portrait photos with *bokeh*—a Japanese term that refers to the pleasing blurring of a background—which photography experts generally consider to be of the highest quality. At that time only expensive single-lens reflex cameras could take such photos, but Hubel thought that with a dual-lens design and advanced computational-photography techniques, Apple could add the capability in the iPhone. His idea aligned well with the camera team's stated purpose: "More people taking better images more of the time."

As the team worked to turn this idea into reality, several challenges emerged. The first attempts produced some amazing portrait pictures but also a number of "failure cases" in which the algorithm was unable to distinguish between the central object in sharp relief (a face, for instance) and the background being blurred. For example, if a person's face was to be photographed from behind chicken wire, it was not possible to construct an algorithm that would capture the chicken wire to the side of the face with the same sharpness as the chicken wire in front of it. The wire to the side would be as blurred as the background.

One might say, “Who cares about the chicken wire case? That’s exceedingly rare.” But for the team, sidestepping rare or extreme situations—what engineers call *corner cases*—would violate Apple’s strict engineering standard of zero “artifacts,” meaning “any undesired or unintended alteration in data introduced in a digital process by an involved technique and/or technology.” Corner cases sparked “many tough discussions” between the camera team and other teams involved, recalls Myra Haggerty, the VP of sensor software and UX prototyping, who oversaw the firmware and algorithm teams. Sebastien Marineau-Mes, the VP to whom the camera software team ultimately reported, decided to defer the release of the feature until the following year to give the team time to better address failure cases—“a hard pill to swallow,” Hubel admits.

To get some agreement on quality standards, the engineering teams invited senior design and marketing leaders to meet, figuring that they would offer a new perspective. The design leaders brought an additional artistic sensibility to the debate, asking, “What makes a beautiful portrait?” To help reassess the zero-artifacts standard, they collected images from great portrait photographers. They noted, among other things, that these

photos often had blurring at the edges of a face but sharpness on the eyes. So they charged the algorithm teams with achieving the same effect. When the teams succeeded, they knew they had an acceptable standard.

Another issue that emerged was the ability to preview a portrait photo with a blurred background. The camera team had designed the feature so that users could see its effect on their photos only *after* they had been taken, but the human interface (HI) design team pushed back, insisting that users should be able to see a "live preview" and get some guidance about how to make adjustments *before* taking the photo. Johnnie Manzari, a member of the HI team, gave the camera team a demo. "When we saw the demo, we realized that this is what we needed to do," Townsend told us. The members of his camera hardware team weren't sure they could do it, but difficulty was not an acceptable excuse for failing to deliver what would clearly be a superior user experience. After months of engineering effort, a key stakeholder, the video engineering team (responsible for the low-level software that controls sensor and camera operations) found a way, and the collaboration paid off. Portrait mode was central to Apple's marketing of the iPhone 7 Plus. It proved a major reason for

users' choosing to buy and delighting in the use of the phone.

As this example shows, Apple's collaborative debate involves people from various functions who disagree, push back, promote or reject ideas, and build on one another's ideas to come up with the best solutions. It requires open-mindedness from senior leaders. It also requires those leaders to inspire, prod, or influence colleagues in other areas to contribute toward achieving their goals.

While Townsend is accountable for how great the camera is, he needed dozens of other teams—each of which had a long list of its own commitments—to contribute their time and effort to the portrait mode project. At Apple that's known as *accountability without control:* You're accountable for making the project succeed even though you don't control all the other teams. This process can be messy yet produce great results. "Good mess" happens when various teams work with a shared purpose, as in the case of the portrait mode project. "Bad mess" occurs when teams push their own agendas ahead of common goals. Those who become associated with bad mess and don't or can't change their behavior are removed from leadership positions, if not from Apple altogether.

Leadership at Scale

Apple's way of organizing has led to tremendous innovation and success over the past two decades. Yet it has not been without challenges, especially with revenues and head count having exploded since 2008.

As the company has grown, entering new markets and moving into new technologies, its functional structure and leadership model have had to evolve. Deciding how to organize areas of expertise to best enable collaboration and rapid decision-making has been an important responsibility of the CEO. The adjustments Tim Cook has implemented in recent years include dividing the hardware function into hardware engineering and hardware technologies; adding artificial intelligence and machine learning as a functional area; and moving human interface out of software to merge it with industrial design, creating an integrated design function.

Another challenge posed by organizational growth is the pressure it imposes on the several hundred VPs and directors below the executive team. If Apple were to cap the size or scope of a senior leader's organization to limit the number and breadth of details that the leader is expected to own, the company would need to hugely

expand the number of senior leaders, making the kind of collaboration that has worked so well impossible to preserve.

Cognizant of this problem, Apple has been quite disciplined about limiting the number of senior positions to minimize how many leaders must be involved in any cross-functional activity. In 2006, the year before the iPhone's launch, the company had some 17,000 employees; by 2019 that number had grown more than eightfold, to 137,000. Meanwhile, the number of VPs approximately doubled, from 50 to 96. The inevitable result is that senior leaders head larger and more diverse teams of experts, meaning more details to oversee and new areas of responsibility that fall outside their core expertise.

In response, many Apple managers over the past five years or so have been evolving the leadership approach described above: experts leading experts, immersion in the details, and collaborative debate. We have codified these adaptions in what we call the *discretionary leadership* model, which we have incorporated into a new educational program for Apple's VPs and directors. Its purpose is to address the challenge of getting this leadership approach to

drive innovation in all areas of the company, not just product development, at an ever-greater scale.

When Apple was smaller, it may have been reasonable to expect leaders to be experts on and immersed in the details of pretty much everything going on in their organizations. However, they now need to exercise greater discretion regarding where and how they spend their time and efforts. They must decide which activities demand their full attention to detail because those activities create the most value for Apple. Some of those will fall within their existing core expertise (what they still need to *own*), and some will require them to *learn* new areas of expertise. Activities that require less attention from the leader can be pushed down to others (and the leaders will either *teach* others or *delegate* in cases where they aren't experts).

Rosner, the VP of applications, provides a good example. Like many other Apple managers, he has had to contend with three challenges arising from Apple's tremendous growth. First, the *size* of his function has exploded over the past decade in terms of both head count (from 150 to about 1,000) and the number of projects under way at any given time. Clearly, he cannot dive into all the details of all those projects. Second, the *scope* of his portfolio has widened:

Over the past 10 years he has assumed responsibility for new applications, including News, Clips (video editing), Books, and Final Cut Pro (advanced video editing). Although apps are his core area of expertise, some aspects of these—among them editorial content for News, how book publishing works, and video editing—involve matters in which Rosner is not an expert. Finally, as Apple's product portfolio and number of projects have expanded, even more coordination with other functions is required, increasing the *complexity* of collaborating across the many units. For instance, whereas Rosner is responsible for the engineering side of News, other managers oversee the operating system on which it depends, the content, and the business relationships with content creators (such as the *New York Times*) and advertisers.

To cope, Rosner has adapted his role. As an expert who leads other experts, he had been immersed in details—especially those concerning the top-level aspects of software applications and their architecture that affect how users engage with the software. He also collaborated with managers across the company in projects that involved those areas.

But with the expansion of his responsibilities, he has moved some things from his *owning* box—including traditional productivity apps such as Keynote and Pages—into his *teaching* box. (See the exhibit "Roger Rosner's discretionary leadership.") Now he guides and gives feedback to other team members so that they can develop software applications according to Apple's norms. Being a teacher doesn't mean that Rosner gives instruction at a whiteboard; rather, he offers strong, often passionate critiques of his team's work. (Clearly, general managers without his core expertise would find it difficult to teach what they don't know.)

Roger Rosner's discretionary leadership

Apple's VP of applications, Roger Rosner, oversees a portfolio comprising four distinct categories that require varying amounts of his time and attention to detail. In 2019 it looked like this:

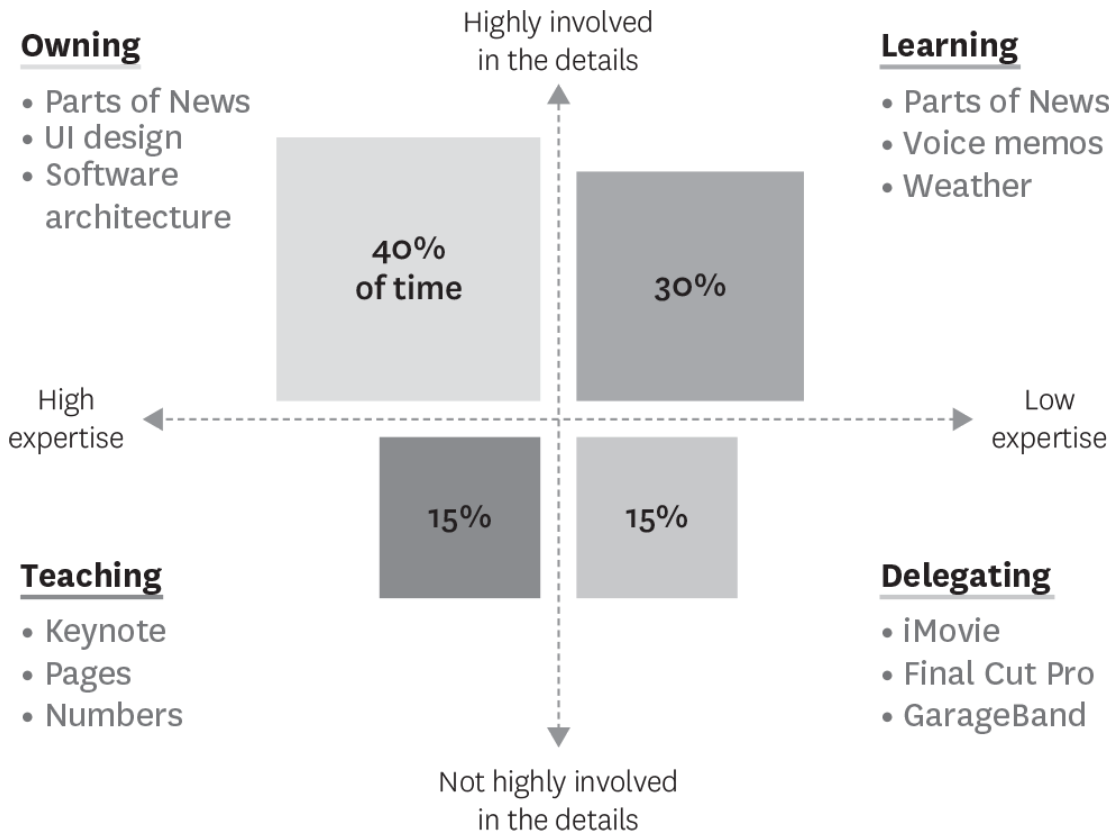

Source: Apple

The second challenge for Rosner involved the addition of activities beyond his original expertise. Six years ago he was given responsibility for the engineering and design of News. Consequently, he had to learn about publishing news content via an app—to understand news publications, digital advertising, machine learning to personalize news content, architecting for privacy, and how to incentivize publishers. Thus some of his work fell into the *learning* box. Here managers face a steep learning

curve to acquire new skills. Given how demanding this is, only critical new activities should fall into this category. Over six years of intense learning, Rosner has mastered some of these areas, which are now in his owning box.

As long as a particular activity remains in the learning box, leaders must adopt a beginner's mindset, questioning subordinates in a way that suggests they don't already know the answer (because they don't). This differs starkly from the way leaders question subordinates about activities in the owning and teaching boxes.

Finally, Rosner has delegated some areas—including iMovie and GarageBand, in which he is not an expert—to people with the requisite capabilities. For activities in the *delegating* box, he assembles teams, agrees on objectives, monitors and reviews progress, and holds the teams accountable: the stuff of general management.

Whereas Apple's VPs spend most of their time in the owning and learning boxes, general managers at other companies tend to spend most of their time in the delegating box. Rosner estimates that he spends about 40% of his time on activities he owns (including collaboration with others in a given area), about 30% on learning, about 15% on teaching, and about 15% on delegating. These

numbers vary by manager, of course, depending on their business and the needs at a given time.

The discretionary leadership model preserves the fundamental principle of an effective functional organization at scale—aligning expertise and decision rights. Apple can effectively move into new areas when leaders like Rosner take on new responsibilities outside their original expertise, and teams can grow in size when leaders teach others their craft and delegate work. We believe that Apple will continue to innovate and prosper by being organized this way.

Apple's functional organization is rare, if not unique, among very large companies. It flies in the face of prevailing management theory that companies should be reorganized into divisions and business units as they become large. But something vital gets lost in a shift to business units: the alignment of decision rights with expertise.

Why do companies so often cling to having general managers in charge of business units? One reason, we believe, is that making the change is difficult. It entails overcoming inertia, reallocating power among managers,

changing an individual-oriented incentive system, and learning new ways of collaborating. That is daunting when a company already faces huge external challenges. An intermediate step may be to cultivate the experts-leading-experts model even within a business unit structure. For example, when filling the next senior management role, pick someone with deep expertise in that area as opposed to someone who might make the best general manager. But a full-fledged transformation requires that leaders also transition to a functional organization. Apple's track record proves that the rewards may justify the risks. Its approach can produce extraordinary results.

Originally published in November–December 2020. Reprint R2006F

About the Contributors

MORRA AARONS-MELE is the founder of the award-winning social impact agency, Women Online, and author of *Hiding in the Bathroom.* She has written for the *New York Times,* the *Wall Street Journal, O the Oprah Magazine,* and other publications, and is the host of *The Anxious Achiever* podcast.

BORIS BABIC is an assistant professor of decision sciences at INSEAD.

HANNAH RILEY BOWLES is the Roy E. Larsen Senior Lecturer in Public Policy and Management at Harvard Kennedy School.

ADAM BRANDENBURGER is the J. P. Valles Professor at the Stern School of Business, a distinguished professor at the Tandon School of Engineering, and the faculty director of the Program on Creativity and Innovation at NYU Shanghai, all at New York University.

TOMAS CHAMORRO-PREMUZIC is the Chief Talent Scientist at ManpowerGroup, a professor of business psychology at University College London and at Columbia

University, and an associate at Harvard's Entrepreneurial Finance Lab. He is the author of *Why Do So Many Incompetent Men Become Leaders? (and How to Fix It)*, upon which his TEDx talk was based. Find him on Twitter: @drtcp.

PRITHWIRAJ (RAJ) CHOUDHURY is the Lumry Family Associate Professor of Business Administration at Harvard Business School. His research focuses on the future of work—especially how work-from-anywhere practices are changing its geography.

I. GLENN COHEN is a deputy dean, professor of law, and faculty director of the Petrie-Flom Center for Health Law Policy, Biotechnology, and Bioethics at Harvard Law School.

ROBIN J. ELY is the Diane Doerge Wilson Professor of Business Administration at Harvard Business School and the faculty chair of the HBS Gender Initiative.

THEODOROS EVGENIOU is a professor of decision sciences and technology management at INSEAD.

FRANCES FREI is the UPS Foundation Professor of Service Management at Harvard Business School. She received shares in Uber as compensation for her work with the company, which she continues to hold. She is one of the authors of *Unleashed: The Unapologetic Leader's Guide to Empowering Everyone Around You* (Harvard Business Review Press, 2020).

GRETCHEN GAVETT is a senior editor at *Harvard Business Review.*

SARA GERKE is a research fellow in medicine, artificial intelligence, and law at the Petrie-Flom Center.

DENNIS A. GIOIA is the Robert and Judith Auritt Klein Professor of Management at the Smeal College of Business at Penn State University.

MORTEN T. HANSEN is a member of Apple University's faculty and a professor at the University of California, Berkeley. He was formerly on the faculties of Harvard Business School and INSEAD.

DOUGLAS HOLT is the founder and president of the Cultural Strategy Group and was formerly a professor

at Harvard Business School and the University of Oxford. He is the author of *How Brands Become Icons: The Principles of Cultural Branding* (Harvard Business School Press, 2004).

ROBERT LIVINGSTON is the author of *The Conversation: How Seeking and Speaking the Truth About Racism Can Radically Transform Individuals and Organizations*. He also serves on the faculty of the Harvard Kennedy School.

ANNE MORRISS is an entrepreneur and the executive founder of the Leadership Consortium. She is one of the authors of *Unleashed: The Unapologetic Leader's Guide to Empowering Everyone Around You* (Harvard Business Review Press, 2020).

BARRY NALEBUFF is the Milton Steinbach Professor of Management at the Yale School of Management, where he teaches negotiation, innovation, strategy, and game theory.

JOEL M. PODOLNY is a vice president of Apple and the dean of Apple University. Prior to joining Apple, in

2009, he was the dean of the Yale School of Management and on the faculty of Harvard's and Stanford's business schools.

DAVID A. THOMAS is the president of Morehouse College. He is also the H. Naylor Fitzhugh Professor Emeritus at Harvard Business School and the former dean of Georgetown University's McDonough School of Business.

BOBBI THOMASON is an assistant professor of applied behavioral science at Pepperdine University.

VERÓNICA H. VILLENA is an assistant professor of supply chain management at the Smeal College of Business at Penn State University.

Index